

# Beyond Airflow Magnitude: Distribution-Aware Machine Learning for Rack-Level Thermal Prediction in IoT-Driven Data Centers

Antonio Cortés Castillo  
Departamento de Informática  
Universidad de Panamá  
Ciudad de Panamá, Panamá  
antonio.cortes@up.ac.pa

**Abstract**— Existing approaches for thermal management in data centers rely predominantly on computational fluid dynamics (CFD) simulations or facility-level machine learning models, which often lack scalability or fail to capture fine-grained airflow–thermal interactions at the rack level. To address this gap, this study proposes a predictive modeling framework that leverages airflow dynamics to estimate rack-level thermal behavior using machine learning techniques. The analysis is conducted on a dataset of 300 observations, incorporating airflow velocity, airflow distribution uniformity, and spatial rack positioning. Unlike conventional approaches that prioritize airflow magnitude, the proposed framework explicitly models airflow distribution as a key driver of thermal performance. Results demonstrate that airflow distribution has a significantly greater impact on maximum temperature and hotspot formation than airflow velocity alone. Among the evaluated models, XGBoost achieved the highest predictive accuracy ( $R^2 = 0.94$ , RMSE = 0.95), effectively capturing non-linear airflow–thermal interactions. The findings provide new insights into the importance of distribution-aware cooling strategies and enable a transition from reactive to predictive thermal management. This work establishes a scalable and data-driven foundation for fine-grained optimization of cooling performance in modern IoT-driven data centers.

**Keywords**— machine learning, airflow distribution, data center cooling, thermal prediction, hotspot detection

## I. INTRODUCTION

The rapid expansion of artificial intelligence (AI), cloud computing, and IoT-driven applications has led to a significant increase in computational density within modern data centers, intensifying thermal management challenges. As power densities rise, traditional cooling approaches struggle to maintain stable operating conditions, leading to higher energy consumption, thermal hotspots, and reduced system reliability. Data centers account for a significant share of global electricity consumption [1], and this demand continues to grow with the expansion of high-performance computing. Despite advances in cooling infrastructure, current strategies largely focus on increasing airflow capacity rather than optimizing airflow behavior, resulting in inefficient thermal management under non-uniform and dynamic conditions. This limitation highlights the need for predictive and distribution-aware approaches capable of capturing complex airflow–thermal interactions at fine-grained levels.

Cooling infrastructure represents a major component of data center energy consumption, as it removes heat generated by IT

equipment. Recent studies indicate that cooling systems can account for 30% to 55% of total energy consumption in data centers, depending on architectural and operational factors [2], [3]. Because nearly all electrical energy consumed by servers is converted to heat, efficient thermal management is critical for cutting operating expenses and advancing overall energy efficiency. As a result, optimizing airflow distribution and heat transfer within data centers has become a significant research focus.

At the rack level, effective airflow management is essential to maintaining stable operating temperatures in servers and networking equipment. Inadequate airflow can lead to recirculation zones, thermal hotspots, and irregular heat gradients, compromising system dependability and hardware lifespan. Recent studies show that improving strategies such as cold aisle containment (CAC) and hot aisle containment (HAC) can greatly improve cooling performance by promoting uniform temperature distribution within data centers [4]. These strategies minimize the mixing of cold air supply with hot server exhaust.

Several physical parameters determine the thermal performance of rack-level cooling systems, including IT load, airflow velocity, pressure drop across server components, and temperature differential between inlet and outlet air streams. Computational and experimental studies show that airflow distribution within racks directly influences heat removal efficiency and temperature variations across servers [5]. Mathematical simulations and computational fluid dynamics (CFD) models are commonly used to characterize airflow patterns and quantify their impact on heat dissipation and temperature distributions inside racks. These approaches provide valuable insights into the operation of cooling systems under various conditions.

Dimensionless parameters provide a solid framework for analyzing convective heat transfer in airflow systems. The Reynolds number (Re) is widely used to characterize airflow regimes and distinguish between laminar and turbulent flow, both of which significantly affect heat transfer performance. The Nusselt number (Nu) quantifies the ratio of convective to conductive heat transfer and serves as a key metric for evaluating heat removal efficiency from surfaces exposed to airflow. These parameters enable researchers to generalize airflow and thermal behavior across a wide range of configurations and operating conditions [6].

In recent years, machine learning techniques have emerged as powerful tools for modeling complex and non-linear relationships in engineering systems. In the context of data center thermal management, these approaches offer the potential to move beyond traditional descriptive and simulation-based analyses toward predictive and adaptive cooling strategies.

Despite extensive research on thermal management in data centers, most existing studies remain limited to descriptive or simulation-based analyses, lacking predictive capabilities that can effectively capture complex airflow–thermal interactions in real-world environments. This limitation restricts the ability to systematically optimize cooling strategies based on airflow behavior.

To address this limitation, this paper proposes a predictive modeling framework that leverages airflow dynamics to estimate thermal performance at the rack level, enabling data-driven and scalable cooling optimization. Furthermore, the study extends these findings to highlight their implications for data center-wide thermal management.

This paper makes the following contributions:

- A quantitative evaluation of airflow dynamics and their impact on rack-level thermal behavior.
- A machine learning-based predictive modeling approach (Random Forest and XGBoost) to estimate thermal performance based on airflow parameters.
- An interpretation of rack-level results in the context of data center-scale cooling efficiency.
- Practical insights for improving airflow management in high-density environments.

The remainder of this paper is structured as follows. Section II reviews related work on airflow optimization and temperature management in data centers. Section III describes the dataset and the proposed methodological framework. Section IV presents the results and discussion derived from both experimental and predictive analyses. Finally, Section V concludes the paper and outlines directions for future research.

## II. RELATED WORK

Thermal management and airflow optimization have become critical research areas due to the increasing computational density and energy consumption of modern data centers. Numerous studies have employed Computational Fluid Dynamics (CFD) and numerical modeling techniques to analyze airflow behavior and temperature distribution within server racks. These approaches provide detailed insights into heat transfer mechanisms and airflow patterns, demonstrating that airflow organization plays a crucial role in heat dissipation and temperature uniformity [8], [9]. However, despite their accuracy, CFD-based methods are computationally intensive and often unsuitable for real-time monitoring or adaptive cooling control.

In addition to simulation-based approaches, several studies have investigated practical airflow management strategies aimed at improving cooling efficiency. Techniques such as

raised-floor air supply, cold aisle containment (CAC), and hot aisle containment (HAC) have been shown to reduce recirculation and improve temperature uniformity within data centers [10], [11]. Experimental and hybrid studies combining physical measurements with numerical simulations further confirm that airflow control mechanisms, including containment systems, can significantly reduce energy consumption and enhance cooling performance [12], [13]. Nevertheless, these approaches primarily focus on static configuration improvements and do not provide predictive capabilities for dynamic thermal behavior.

More recently, machine learning (ML) techniques have emerged as promising tools for modeling complex and non-linear relationships in data center environments. Surrogate modeling approaches enable the approximation of airflow and thermal dynamics with significantly reduced computational cost compared to traditional CFD simulations. These models facilitate faster prediction of temperature fields and airflow behavior, enabling more efficient monitoring and control of data center operations [14], [15]. Despite these advancements, most ML-based studies focus on facility-level optimization and lack fine-grained analysis at the rack level, where localized airflow inefficiencies and hotspot formation occur.

In parallel, advanced optimization techniques, including physics-informed machine learning and deep learning-based approaches, have been proposed to enhance cooling system performance. These methods enable adaptive control strategies that improve energy efficiency while maintaining safe operating conditions [16], [17]. Although promising, such approaches often require large-scale data and may not explicitly capture the detailed airflow–thermal interactions within individual racks.

The increasing adoption of high-density computing infrastructures, including AI and large-scale data analytics systems, has further intensified thermal challenges in modern data centers. Recent studies highlight that rising power densities introduce non-linear airflow patterns and complex thermal dynamics, making traditional cooling strategies less effective [18]. To address these challenges, deep learning-based models have been explored for predicting thermal behavior and identifying hotspot formation with higher temporal resolution compared to conventional simulation methods [19]. Furthermore, optimization techniques leveraging advanced data-driven models have been proposed to generate efficient airflow configurations and improve cooling performance under dynamic conditions [20].

In addition to airflow dynamics, the physical layout and network topology of data centers also influence thermal behavior. Architectural designs such as Fat-Tree and BCube affect component distribution, cabling density, and airflow pathways within racks. These structural factors can create airflow obstructions, increase pressure drop, and lead to localized heat accumulation [21]. Consequently, accurate thermal modeling must consider not only airflow parameters but also infrastructure-level constraints that impact cooling efficiency.

Despite these advances, a significant gap remains in developing predictive, data-driven models that accurately capture airflow–thermal interactions at the rack level while maintaining scalability for data center-wide optimization. In contrast to existing approaches, this study integrates machine learning techniques with experimental airflow analysis to establish predictive relationships between airflow distribution and thermal performance. This approach enables more accurate, scalable, and actionable insights for optimizing cooling strategies in modern data center environments.

### III. METHODOLOGY

This study employs a hybrid methodology combining experimental data analysis, dimensionless physical modeling, and machine learning techniques to investigate airflow-driven thermal behavior at the rack level. This approach enables both physical interpretability and predictive capability, addressing limitations of purely statistical or simulation-based methods.

#### A. Experimental Dataset Description

The dataset used in this study comprises 300 observations collected from a rack-level cooling environment designed to analyze airflow-driven thermal behavior in data centers. Each observation captures key airflow characteristics, spatial conditions, and corresponding thermal responses within the rack. The dataset is available at: URL:[https://github.com/acortescastillo/dataset\\_airflow\\_thermal-300.git](https://github.com/acortescastillo/dataset_airflow_thermal-300.git).

The dataset includes five primary variables: airflow velocity (`airflow_mps`), airflow distribution index (`distribution_idx`), rack position (`rack_level`), maximum temperature (`max_temp_c`), and hotspot occurrence (`hotspots`). Airflow velocity represents the magnitude of the cooling flow, while the distribution index quantifies the uniformity of airflow across the rack. The `rack_level` variable encodes the vertical position within the rack (e.g., bottom, middle, top), enabling spatial analysis of thermal gradients.

The thermal response is characterized using two complementary indicators: the maximum rack temperature (`max_temp_c`), which captures overall thermal conditions, and hotspot occurrence (`hotspots`), which reflects localized thermal anomalies. This dual representation enables a comprehensive assessment of both global and localized cooling performance.

All variables are continuous or ordinal and were validated to ensure completeness and consistency, with no missing values present in the dataset. This guarantees reliable statistical analysis and supports the development of robust predictive models.

The selected variables provide a compact yet expressive representation of airflow–thermal interactions, enabling both interpretable statistical analysis and effective machine learning-based prediction of rack-level thermal behavior.

Additionally, the dataset was designed to capture both uniform and non-uniform airflow conditions, enabling the analysis of their impact on thermal performance and hotspot formation.

#### B. Physical and Dimensionless Analysis

To provide a physically grounded interpretation of airflow and heat transfer behavior, dimensionless parameters were incorporated as supporting analytical tools. The Reynolds number ( $Re$ ) was used to characterize airflow regimes and distinguishing between laminar and turbulent flow conditions, while the Nusselt number ( $Nu$ ) was used to evaluate convective heat transfer efficiency.

These parameters were analyzed to complement the primary airflow variables, particularly airflow velocity (`airflow_mps`), by providing insight into the underlying fluid dynamics governing heat dissipation.

The relationship between Reynolds and Nusselt numbers was modeled using the classical power-law formulation:

$$Nu = C \cdot Re^m \quad (1)$$

Which describes the dependence of convective heat transfer on flow conditions. To facilitate regression analysis, a logarithmic transformation was applied:

$$\log(Nu) = \log(C) + m \log(Re) \quad (2)$$

This analysis provides a theoretical foundation for understanding airflow-thermal interactions, supporting the interpretation of experimental observations and the development of predictive models.

While these dimensionless parameters are not directly used as primary predictors in the machine learning models, they offer valuable physical insight into the mechanisms driving thermal behavior within the rack.

#### C. Statistical and Regression Analysis

To analyze the relationship between airflow characteristics and thermal behavior at the rack level, a statistical and regression-based framework was employed using the variables airflow velocity (`airflow_mps`), airflow distribution index (`distribution_idx`), rack position (`rack_level`), maximum temperature (`max_temp_c`), and hotspot occurrence (`hotspots`).

Initially, descriptive statistical analysis was conducted to examine the distribution and variability of the selected variables. Correlation analysis was then performed to quantify the relationships between airflow parameters and thermal response indicators, enabling the identification of dominant factors influencing temperature behavior and hotspot formation.

Additionally, hotspot occurrence was modeled as a discrete indicator to capture critical thermal conditions, complementing continuous temperature-based analysis.

To model the relationship between airflow characteristics and thermal performance, a multivariable regression model was defined, where the maximum rack temperature serves as the primary dependent variable:

$$T_{max} = \beta_0 + \beta_1 \cdot \text{airflow\_mps} + \beta_2 \cdot \text{distribution\_idx} + \beta_3 \cdot \text{rack\_level} \quad (3)$$

where  $T_{max}$  represents the maximum temperature within the rack, and  $\beta_0, \beta_1, \beta_2,$  and  $\beta_3$  are regression coefficients. This model captures the linear contribution of airflow magnitude, distribution uniformity, and spatial position to thermal behavior.

In addition to temperature modeling, hotspot occurrence was analyzed as a secondary response variable to evaluate localized thermal anomalies. This analysis enables a more comprehensive understanding of how airflow distribution affects not only overall temperature levels but also the formation of critical thermal regions.

Scatter plot analysis and residual evaluation were further conducted to identify potential non-linear relationships and deviations from linear assumptions, providing justification for the subsequent application of machine learning models.

#### D. Predictive Modeling Framework

To capture the complex and non-linear relationships between airflow characteristics and thermal behavior, a predictive modeling framework based on supervised machine learning was developed. The framework leverages key airflow and spatial variables—airflow velocity (`airflow_mps`), airflow distribution index (`distribution_idx`), and rack position (`rack_level`)—to estimate thermal response indicators, including maximum temperature (`max_temp_c`) and hotspot occurrence (`hotspots`).

Unlike conventional regression approaches, which primarily capture linear dependencies, the proposed framework is designed to model higher-order interactions between airflow magnitude, distribution uniformity, and spatial positioning within the rack. These interactions are critical for accurately representing localized thermal phenomena and non-uniform cooling behavior.

Two ensemble learning methods were employed: Random Forest (RF) and Extreme Gradient Boosting (XGBoost). These models were selected due to their robustness in handling structured data, their ability to capture non-linear feature interactions, and their strong predictive performance in engineering applications. In addition, a linear regression model was used as a baseline to provide a reference for evaluating performance improvements achieved through non-linear modeling.

The predictive task was formulated in two complementary ways. First, a regression approach was used to estimate the continuous thermal response variable (`max_temp_c`). Second, hotspot occurrence (`hotspots`) was modeled to capture localized thermal anomalies, enabling a more comprehensive understanding of airflow-driven thermal behavior.

The dataset was partitioned into training and testing subsets using an 80/20 split to ensure robust model validation. Additionally, model stability was verified across multiple random splits to ensure robustness of the predictive performance. To reduce overfitting, models were evaluated using train-test separation and performance consistency analysis. Model performance was assessed using the coefficient of determination ( $R^2$ ) and Root Mean Square Error (RMSE) for temperature prediction, while hotspot-related predictions were analyzed to evaluate the model's ability to capture critical thermal conditions.

Furthermore, feature importance analysis was conducted to quantify the relative contribution of airflow velocity, distribution uniformity, and rack position in determining

thermal outcomes. This analysis provides interpretability to the predictive models and supports the identification of dominant factors influencing cooling performance.

Overall, this predictive framework enables accurate estimation of thermal behavior under varying airflow conditions, supporting data-driven optimization of cooling strategies and addressing limitations of traditional descriptive and simulation-based approaches.

This framework enables the transition from reactive cooling strategies to predictive and adaptive thermal management.

#### E. System-Level Implications

Although the analysis is conducted at the rack level, the findings have direct implications for data center-wide thermal management. Airflow velocity, distribution uniformity, and spatial positioning within racks significantly influence both global temperature levels and localized hotspot formation. When these effects are not properly managed, localized inefficiencies can propagate across the facility, leading to reduced cooling effectiveness and increased energy consumption.

The ability to accurately estimate maximum temperature and hotspot occurrence using airflow-related variables enables a more proactive approach to thermal management. Instead of relying on reactive cooling strategies, predictive insights can be used to anticipate thermal imbalances and adjust airflow distribution dynamically. This is particularly relevant in high-density environments, where small variations in airflow behavior can lead to significant thermal deviations.

From an operational perspective, these insights support the development of intelligent cooling strategies, including adaptive airflow control, optimized workload placement, and improved rack-level configuration. By prioritizing airflow distribution alongside airflow magnitude, it becomes possible to enhance cooling efficiency without proportionally increasing energy consumption.

Furthermore, the integration of predictive models into data center management systems enables scalable optimization across multiple racks and zones. This facilitates coordinated thermal control, reducing the likelihood of hotspot propagation and improving overall system reliability.

Ultimately, the results highlight the importance of transitioning toward data-driven and distribution-aware cooling strategies, where airflow behavior is continuously monitored and optimized to achieve both thermal stability and energy efficiency in modern data center infrastructures.

## IV. RESULT AND DISCUSSION

The results demonstrate that airflow characteristics—particularly airflow velocity, distribution uniformity, and spatial positioning—have a significant impact on rack-level thermal behavior. The analysis reveals that both maximum temperature (`max_temp_c`) and hotspot occurrence (`hotspots`) are strongly influenced by how airflow is distributed within the rack, rather than by airflow magnitude alone.

Experimental observations indicate that configurations with improved airflow distribution (higher `distribution_idx` values)

consistently exhibit lower maximum temperatures and reduced hotspot formation. In contrast, scenarios with high airflow velocity but poor distribution show limited thermal improvement, confirming that non-uniform airflow leads to localized overheating despite increased cooling capacity.

A clear non-linear relationship was identified between airflow velocity (airflow\_mps) and thermal response. As airflow increases, maximum temperature decreases initially; however, beyond a threshold of approximately 3.0 m/s, the rate of temperature reduction diminishes significantly. This behavior highlights the presence of diminishing returns and demonstrates that increasing airflow alone is insufficient for achieving optimal cooling performance.

The influence of spatial positioning (rack\_level) further reveals the presence of vertical thermal gradients within the rack. Upper rack levels tend to exhibit higher temperatures and a greater concentration of hotspots due to heat accumulation and airflow recirculation effects. This finding underscores the importance of considering spatial airflow distribution when designing cooling strategies.

The predictive modeling results reinforce these observations by capturing complex interactions between airflow variables and thermal behavior. As summarized in Table I, the XGBoost model achieved the highest predictive performance, with an  $R^2$  value of 0.94 and an RMSE of 0.95, outperforming Random Forest ( $R^2 = 0.91$ ) and linear regression ( $R^2 = 0.82$ ).

TABLE I. PERFORMANCE COMPARISON OF PREDICTIVE MODELS FOR RACK-LEVEL THERMAL ESTIMATION

Model	$R^2$	RMSE
Linear Regression	0.82	1.45
Random Forest	0.91	1.1
XGBoost	0.94	0.95

These results demonstrate the effectiveness of machine learning in modeling non-linear airflow-thermal relationships that cannot be fully captured using traditional regression techniques.

Fig. 1 illustrates the relationship between airflow velocity and maximum temperature, highlighting the diminishing returns in cooling performance at higher airflow rates.

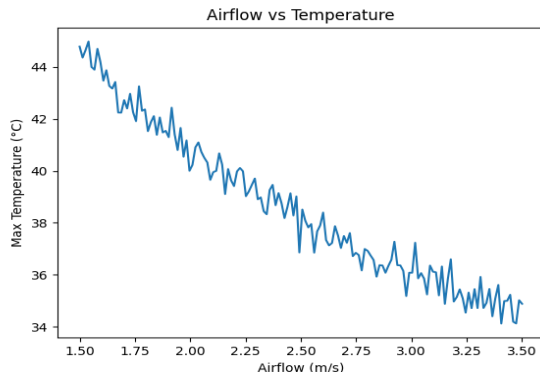


Fig. 1. Relationship between airflow velocity and maximum rack temperature.

Fig. 2 presents the comparison between predicted and actual temperature values, confirming the accuracy and reliability of the XGBoost model.




Fig. 2. Comparison between predicted and actual maximum temperature using the XGBoost model.

Additionally, Fig. 3 shows the spatial distribution of temperature within the rack, where localized hotspots are clearly observed under non-uniform airflow conditions.

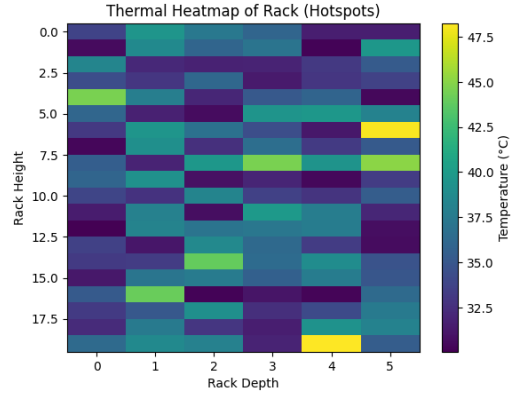


Fig. 3. Thermal distribution within the rack showing hotspot formation under non-uniform airflow conditions.

Feature importance analysis indicates that airflow distribution (distribution\_idx) is the most influential factor affecting thermal performance, followed by airflow velocity (airflow\_mps) and rack position (rack\_level). This result provides strong evidence that optimizing airflow distribution is more critical than simply increasing airflow magnitude.

From a practical standpoint, these findings suggest that effective cooling strategies should prioritize airflow uniformity and spatial optimization to reduce thermal gradients and prevent hotspot formation. The integration of predictive models enables data-driven decision-making, allowing operators to anticipate thermal issues and implement proactive cooling adjustments.

Overall, the results validate the proposed approach and demonstrate its capability to provide both accurate predictions and actionable insights, addressing key limitations of traditional

descriptive and simulation-based studies in data center thermal management.

## V. CONCLUSION

This study investigated the relationship between airflow characteristics and thermal behavior at the rack level, focusing on airflow velocity, distribution uniformity, and spatial positioning. The results demonstrate that airflow distribution significantly influences maximum temperature and hotspot formation within the rack, confirming that thermal performance depends more on airflow uniformity than on airflow magnitude alone.

The analysis revealed the presence of non-linear airflow-thermal interactions, where increasing airflow velocity leads to diminishing improvements in cooling efficiency beyond a certain threshold. Additionally, spatial variations within the rack contribute to vertical thermal gradients, with upper regions more prone to heat accumulation and hotspot formation.

The proposed predictive modeling framework, based on machine learning techniques, successfully captured these complex relationships. In particular, the XGBoost model achieved the highest predictive accuracy, demonstrating its effectiveness in modeling non-linear thermal behavior and providing reliable temperature estimation under varying airflow conditions. The inclusion of hotspot occurrence as a complementary indicator further enhanced the understanding of localized thermal phenomena.

From a practical perspective, the findings highlight the importance of prioritizing airflow distribution and spatial optimization in cooling system design. Rather than relying solely on increased airflow capacity, data-driven strategies that account for airflow uniformity and rack-level variability can significantly improve cooling efficiency and reduce energy consumption.

Furthermore, the ability to predict thermal behavior using airflow-related variables enables the development of proactive and adaptive cooling strategies. These approaches can be extended to data center-scale operations, supporting intelligent workload placement, coordinated airflow management, and improved overall system reliability.

Although the dataset size is moderate (300 observations), the results demonstrate consistent trends across different airflow conditions, supporting the validity of the proposed approach and its applicability to real-world data center environments.

These findings position airflow-aware predictive modeling as a key enabler for intelligent, energy-efficient, and scalable data center infrastructures.

Future work will focus on expanding the dataset to include additional operational conditions and exploring real-time integration of predictive models into data center management systems. The incorporation of advanced learning techniques and larger-scale validation scenarios will further enhance the applicability and robustness of the proposed approach.

## REFERENCES

- [1] A. Shehabi et al., "United States data center energy usage report," Lawrence Berkeley National Laboratory, 2020.
- [2] H. Wang, J. Huang, and L. Zhang, "Data centers cooling: A critical review of techniques, challenges, and energy saving solutions," *International Journal of Refrigeration*, 2024.
- [3] N. Li et al., "Evaluation of the cooling effectiveness of air-cooled data centers using energy diagrams," *Applied Energy*, vol. 382, 2025.
- [4] D. Jiang, "Effects and optimization of airflow on the thermal environment in a data center," *Frontiers in Built Environment*, vol. 10, 2024.
- [5] S. Zhang, Y. Liu, and X. Chen, "Numerical study of fluid flow and temperature distributions in a data center," *Case Studies in Thermal Engineering*, 2021.
- [6] T. Bergman, F. Incropera, D. DeWitt, and A. Lavine, *Fundamentals of Heat and Mass Transfer*, 8th ed., Wiley, 2020.
- [7] X. Zhan et al., "Data center cooling system optimization using offline reinforcement learning," 2025.
- [8] T.-F. Yang, Y.-C. Chen, B.-L. Chen, C.-H. Li, and W.-M. Yan, "Numerical study of fluid flow and temperature distributions in a data center," *Case Studies in Thermal Engineering*, vol. 28, p. 101405, 2021, doi: 10.1016/j.csite.2021.101405.
- [9] A. M. N. Chacón, F. B. S. Oliveira, and D. S. Dominguez, "Modeling and Thermal Simulation of 2U-Rack Servers Applying Computational Fluid Dynamics Techniques," *DDF*, vol. 427, pp. 121–133, Jul. 2023, doi: 10.4028/p-mid78f.
- [10] Zhao, J.; Wu, J.; Li, M. Experimental and Numerical Investigation of Airflow Organization in Modular Data Centres Utilizing Floor Grid Air Supply. *Buildings* 2024, 14, 2750. <https://doi.org/10.3390/buildings14092750>
- [11] B. Ma, "Simulation and experimental research on the optimization of airflow organization and energy saving in data centers using air deflectors", *Clean Energy Sci. Technol.*, vol. 2, no. 1, p. 141, Mar. 2024.
- [12] Wang, F.; Huang, Y.; Prasetyo, B. Energy-Efficient Improvement Approaches through Numerical Simulation and Field Measurement for a Data Center. *Energies* 2019, 12, 2757. <https://doi.org/10.3390/en12142757>
- [13] J. Choi, J. Hong, T. Hong, S. Jo, J. Kim, and H. Jang, "Numerical analysis of the individual and combined effects of the aisle containment layouts on the cooling performance of the edge data center," *Building and Environment*, vol. 287, pt. A, Art. no. 113770, 2026, doi: 10.1016/j.buildenv.2025.113770.
- [14] Sarkar, S., Antonio Guillen-Perez, Zachariah Carmichael, Avisek Naug, Refik Mert Çam, Vineet Gundecha, Ashwin Ramesh Babu, Sahand Ghorbanpour and Ricardo Luna Gutiérrez. "Fast 3D Surrogate Modeling for Data Center Thermal Management." *ArXiv abs/2511.11722* (2025): n. pag. 10. [https://arxiv.org/abs/2511.11722?utm\\_source=chatgpt.com](https://arxiv.org/abs/2511.11722?utm_source=chatgpt.com)
- [15] J. Lin, W. Lin, W. Lin, J. Wang, and H. Jiang, "Thermal prediction for air-cooled data center using data-driven-based model," *Applied Thermal Engineering*, vol. 217, Art. no. 119207, 2022, doi: 10.1016/j.applthermaleng.2022.119207.
- [16] X. Zhan, X. Zhu, P. Cheng, X. Hu, Z. He, H. Geng, J. Leng, H. Zheng, C. Liu, T. Hong, Y. Liang, Y. Liu, and F. Zhao, "Data center cooling system optimization using offline reinforcement learning," *arXiv preprint*, arXiv:2501.15085, 2025. [Online]. Available: <https://arxiv.org/abs/2501.15085>
- [17] D. Chen, C.-K. Chui, and P. S. Lee, "Physics informed machine learning based predictive control for intelligent operation of edge datacenters," *Applied Energy*, vol. 402, part B, p. 126975, 2026, doi: 10.1016/j.apenergy.2025.126975
- [18] A. C. Castillo, "Data Center Infrastructure in the Era of Generative Artificial Intelligence," 2025 5th International Conference on Innovative Research in Applied Science, Engineering and Technology (IRASET), Fez, Morocco, 2025, pp. 1-9, doi: 10.1109/IRASET64571.2025.11008109.
- [19] A. C. Castillo, "Application of Generative Artificial Intelligence from Neural Networks in Data Centers," 2025 7th International Congress on Human-Computer Interaction, Optimization and Robotic Applications (ICHORA), Ankara, Turkiye, 2025, pp. 1-8, doi: 10.1109/ICHORA65333.2025.11017008.
- [20] A. C. Castillo, "Neural Network Architectures for Enhancing Data Center Intelligence through Deep Learning," 2025 9th International Symposium on Multidisciplinary Studies and Innovative Technologies (ISMSIT), Ankara, Turkiye, 2025, pp. 1-7, doi: 10.1109/ISMSIT67332.2025.11268124.  
A. C. Castillo, "Various Network Topologies and an Analysis Comparative Between Fat-Tree and BCube for a Data Center Network: An Overview," *2022 IEEE Cloud Summit*, Fairfax, VA, USA, 2022, pp. 1-8, doi: 10.1109/CloudSummit54781.2022.00007.