Journal of Industrial Engineering and Applied Science
Vol. 4, No. 1, 2026 | ISSN 3005-6071 (Print) | ISSN 3005-608X (Online) | CODEN: JIEAAE

Carbon-Emission Estimation Models: Hierarchical Measurement From Board to Datacenter

LIU, Wenwen [1*]

[1] Bytedance, CN

**** LIU, Wenwen is the corresponding author, E-mail: liuwenwen.jessica@bytedance.com***

Abstract: Data centers have become a core contributor to global digital carbon emissions, with their carbon footprint growing 19% annually alongside the expansion of AI and cloud services. Traditional carbon accounting methods are either trapped in macro-level rough calculation based on Power Usage Effectiveness (PUE) or limited to micro-level hardware power consumption measurement, failing to establish a traceable correlation between chip-level energy behavior and datacenter-wide carbon emissions. To address this gap, this study proposes a Hierarchical Coupling Carbon Emission Estimation Model (HCCEEM) that integrates physical modeling and graph neural network (GNN)-based statistical aggregation. The model constructs a four-level traceability chain spanning board (chip), server node, rack cluster, and campus datacenter, and introduces a real-time load adaptation module to capture dynamic workload impacts. Validated on a 14-month dataset from a heterogeneous cloud datacenter, HCCEEM achieves an estimation accuracy of 95.7%, reducing mean absolute error (MAE) by 27.1% and 19.3% compared to PUE-based models and single-level machine learning models respectively. Moreover, the model realizes fine-grained attribution of carbon contributions across levels, revealing that chip-level dynamic power consumption drives 65.2% of server emissions, and rack-level cooling losses account for 33.8% of datacenter emissions. This research provides an interpretable, scalable tool for targeted carbon reduction, bridging the gap between hardware-level optimization and datacenter-wide carbon management. Specifically, HCCEEM exhibits remarkable applicability in high-load scenarios such as large language model (LLM) training and inference, where it can reduce carbon accounting errors by over 30% compared to conventional methods. For small and medium-sized datacenters with limited monitoring resources, the model's modular design allows lightweight deployment by simplifying partial hierarchical modules without significant accuracy loss. Additionally, the hierarchical contribution quantification function of HCCEEM can directly support enterprises' carbon disclosure and compliance reporting, aligning with the carbon neutrality requirements of the digital industry in various regions.

Keywords: Carbon Emission Estimation, Hierarchical Traceability, Cross-Level Coupling, Graph Neural Network, Green Datacenter.

Disciplines: Computer Science. **Subjects:** Artificial Intelligence.

DOI: https://doi.org/10.70393/6a69656173.333931 **ARK:** https://n2t.net/ark:/40704/JIEAS.v4n1a06

1 INTRODUCTION

The global digital economy's rapid growth has pushed datacenter electricity consumption to 290 TWh in 2024, accounting for 1.9% of total global electricity use and generating 152 million tons of carbon emissions. This trend is exacerbated by the proliferation of high-performance computing (HPC) facilities for large language model (LLM) training and inference, where a single 175B-parameter model training process can emit carbon equivalent to 500 cars' annual emissions. According to the Digitalization and Energy Report 2024 released by the International Energy Agency (IEA), the energy consumption of AI-related datacenters will grow at a compound annual growth rate (CAGR) of 25% from 2024 to 2030, far exceeding the average growth rate of 6% for conventional datacenters. This explosive growth has posed severe challenges to the carbon neutrality goals of countries worldwide, especially for cloud service providers and internet enterprises that operate large-scale datacenter clusters. Meanwhile, the heterogeneity of datacenter infrastructure—such as the coexistence of GPU, CPU, and ARM-based servers—and the volatility of workloads—such as sudden spikes in LLM inference queries during promotional events—have further increased the difficulty of accurate carbon emission estimation.

Against this backdrop, accurate carbon-emission estimation is a prerequisite for effective carbon reduction, but current practices face two intractable problems. First, there is a hierarchical decoupling gap: macro-level models rely on PUE and total electricity consumption to calculate emissions, ignoring hardware heterogeneity (e.g., GPU servers consume 4x more energy than CPU-only servers) and dynamic load

fluctuations, leading to estimation errors of up to 35%. Micro-level models, on the other hand, focus on chip power consumption via tools like Intel RAPL, but fail to account for energy losses in power distribution, cooling systems, and network transmission when scaling up to rack or datacenter levels. Second, there is a model interpretability dilemma: existing data-driven models (e.g., random forests, neural networks) achieve high prediction accuracy but act as "black boxes," unable to explain how chip-level parameter adjustments affect datacenter-wide emissions, thus providing limited guidance for targeted optimization.

To solve these problems, this study designs a hierarchical coupling estimation framework with three core innovations. First, it establishes a traceable four-level carbon transmission chain, clarifying the quantitative relationship between each level's energy consumption and emissions. Second, it develops a hybrid modeling engine combining physics-based chip power calculation and GNN-based cross-level aggregation, balancing accuracy and interpretability. Third, it integrates a real-time load adaptation module to handle extreme workload fluctuations (e.g., LLM inference QPS spikes). The framework is validated on a real-world heterogeneous datacenter dataset, providing actionable insights for both hardware engineers and datacenter operators.

2 CRITICAL REVIEW OF EXISTING METHODS

This section systematically reviews three mainstream carbon-emission estimation methods, analyzing their inherent limitations to highlight the necessity of hierarchical coupling modeling.

2.1 MACRO-LEVEL PUE-DRIVEN ROUGH CALCULATION

The PUE-based method is the most widely used industry standard, with the core formula:

Carbon Emissions=Total Facility Energy × Grid Carbon Intensity.

The Green Grid defines PUE as the ratio of total datacenter energy consumption to IT equipment energy consumption, which reflects the energy efficiency of non-IT systems such as cooling and power distribution. Studies by Zhang et al. (2024) and Li et al. (2023) have optimized this method by introducing dynamic PUE correction factors, reducing estimation errors to 15–20% under stable load conditions. However, this method has two fatal flaws: it treats all IT equipment as homogeneous, ignoring the huge energy consumption gap between GPU and CPU servers; and it uses static PUE values, which cannot reflect real-time load changes—for example, PUE increases by 0.1–0.2 during peak LLM inference periods, leading to significant estimation deviations..

2.2 MICRO-LEVEL HARDWARE POWER CONSUMPTION MEASUREMENT

Micro-level methods focus on chip or server power consumption, using physical formulas or sensor data to calculate energy use. The classic CMOS dynamic power consumption formula $P = CV^2f + P_{static}$ is the foundation of chip-level modeling, where C is capacitance, V is voltage, f is frequency, and P_{static} is static leakage power. Tools like Intel RAPL and NVIDIA NVML can measure CPU/GPU power consumption in real time with an accuracy of 90–95%. Wang et al. (2023) proposed a dynamic power correction model based on task type, improving chip-level estimation accuracy to 96%. However, these methods cannot scale to higher levels, as they ignore energy losses in power distribution (e.g., PDU efficiency of 92% means 8% energy loss) and cooling systems (e.g., a 300W server requires an additional 150W for cooling). This leads to a 50% underestimation of emissions when directly extrapolating chip-level data to datacenter level.Another critical limitation is the hardware architecture dependency of micro-level methods: for example, ARM-based servers have lower static leakage power than x86-based servers, but existing physical models often use parameters calibrated for x86 chips, resulting in 10–15% estimation errors for ARM clusters. Furthermore, deploying high-precision sensors for chip-level power measurement requires substantial capital investment, which is unaffordable for small and medium-sized datacenters with limited budgets.

2.3 DATA-DRIVEN BLACK BOX PREDICTION

In recent years, machine learning models have been widely used for carbon-emission estimation. Chen et al. (2024) used a random forest model to predict datacenter emissions based on workload, temperature, and PUE data, achieving an accuracy of 89.2%. Liu et al. (2023) proposed a deep neural network model that integrates hardware utilization and energy supply data, reducing MAE by 14% compared to traditional methods. However, these models have two critical limitations: first, they lack hierarchical structure, treating the datacenter as a single black box and unable to trace emission sources; second, they rely on large amounts of labeled data, which is difficult to collect for small and medium-sized datacenters. More importantly, their black-box nature makes it impossible to explain the causal relationship between input factors and emission outcomes, limiting their practical application for carbon reduction decision-making.The poor generalization ability of data-driven models is another prominent problem: a model trained on a cloud datacenter dataset may only achieve 60–70% accuracy when applied to an edge datacenter, due to differences in infrastructure scale, cooling methods, and workload characteristics. Additionally, the training process of these models is often time-consuming and computationally intensive — training a deep neural network model for a medium-sized datacenter may require hundreds of GPU hours,

increasing the carbon footprint of the model itself.

3 HIERARCHICAL COUPLING CARBON EMISSION ESTIMATION FRAMEWORK

This section presents the core design of HCCEEM, including the four-level traceability logic, cross-level coupling modeling engine, and real-time adaptation mechanism.

3.1 FOUR-LEVEL CARBON EMISSION TRACEABILITY LOGIC

The framework constructs a top-down emission transmission chain spanning board (chip) → server node → rack cluster → campus datacenter, where each level's carbon emissions are determined by the lower-level emissions plus the energy losses of the current level's infrastructure. The core traceability rules are as follows:

1.Board Level: Chip dynamic power consumption is the core emission source, determined by hardware parameters (frequency, voltage) and dynamic load (CPU/GPU utilization, task type). Emissions at this level are calculated based on physical formulas, with no infrastructure losses.

2.Server Node Level: Aggregates chip-level power consumption of all components (CPU, GPU, memory, storage) and adds energy losses from the power supply unit (PSU). The PSU efficiency factor η_{PSU} ,is introduced to quantify losses, with typical values ranging from 88% to 96%.

3.Rack Cluster Level: Aggregates emissions from all servers in the rack, plus losses from power distribution units (PDU) and rack-level cooling systems. The PDU efficiency factor η_{PDU} and cooling loss factor λ_{cool} are used to measure energy waste, with λ_{cool} increasing with server density.

4.Campus Datacenter Level: Aggregates emissions from all racks, plus losses from central cooling systems, network equipment, and auxiliary facilities. The final emissions are amplified by the campus energy structure factor (grid carbon intensity α_{grid} and renewable energy ratio β_{renew}).

To illustrate the application of this traceability chain, we take a single GPU server as an example: the chip-level power consumption of an NVIDIA A100 GPU is 400W under full load, with a CPU power consumption of 150W; at the server node level, the PSU efficiency is 92%, so the total server power consumption is (400+150)/0.92≈598W; at the rack cluster level, the PDU efficiency is 95% and the cooling loss factor is 0.3, so the rack-level power consumption for this server is 598/0.95×(1+0.3)≈823W; at the campus datacenter level, the grid carbon intensity is 0.58 kg CO_2/kWh and the renewable energy ratio is 20%, so the final carbon emissions of this server are $823\times10^{-3}\times0.58\times(1-0.2)$=0.385kg CO_2 per hour. This example clearly shows how emissions propagate across levels, providing a transparent basis for emission source tracing.

3.2 CROSS-LEVEL COUPLING MODELING ENGINE

The modeling engine integrates two complementary modules to balance accuracy and interpretability.

1.Chip-Level Physics-Based Power Calculation Module: For the board level, the classic CMOS dynamic power formula is modified to incorporate dynamic load factors:

$$P_{chip}=(C\times V^2\times f\times U+P_{static}$$

Where U is the CPU/GPU utilization rate (0–100%), which dynamically adjusts power consumption based on real-time workload. P_{static} is calibrated using Intel RAPL data, accounting for leakage power at idle state.Chip-level carbon emissions are calculated as $CE_{chip}=P_{chip}\times\alpha_{grid}$.

2.GNN-Based Cross-Level Aggregation Module: For server, rack, and datacenter levels, a graph neural network is used to aggregate lower-level emissions and quantify infrastructure losses. Each level is treated as a node in the graph, and edges represent emission transmission relationships. The GNN model takes level-specific factors (e.g., PSU efficiency for servers, PDU efficiency for racks) as node features, and learns the emission aggregation rules from historical data. Compared to traditional tree-based models, GNN better captures the hierarchical dependencies between levels, improving aggregation accuracy by 12–15%.

The GNN model used in this study adopts a two-layer Graph Convolutional Network (GCN) structure, with each node feature vector dimension set to 12, including hardware parameters, efficiency factors, and load data. The ReLU activation function is used to introduce non-linearity, and the Adam optimizer is employed to minimize the mean squared error between predicted and actual emissions. The integration of the physical model and GNN follows a sequential workflow: first, the physical model calculates chip-level power consumption with high interpretability; then, the GNN model aggregates the chip-level results to higher levels, capturing complex infrastructure loss relationships that are difficult to describe with physical formulas. This hybrid design avoids the black-box problem of pure data-driven models while maintaining high estimation accuracy.

3.3 REAL-TIME LOAD ADAPTATION MECHANISM

To handle dynamic workload fluctuations (e.g., LLM inference QPS spikes from 2k to 12k), a sliding window adaptation mechanism is designed. The mechanism updates the GNN model parameters every 5 minutes using the latest 24 hours of load data, ensuring that the model can adapt to sudden changes in task type and utilization rate. The adaptation process includes three steps: (1) real-time collection of load data (CPU/GPU utilization, sequence length for LLM tasks); (2) calculation of load fluctuation

coefficients to identify peak periods; (3) retraining the GNN model with the latest data window to adjust aggregation weights. This mechanism reduces estimation error by 18.7% during peak load periods compared to static models. The selection of the sliding window parameters is based on a comprehensive analysis of workload characteristics: the 5-minute update interval is determined by the response time requirement of LLM inference tasks, while the 24-hour data window ensures that the model captures both short-term load spikes and long-term load trends. To further reduce computational overhead, a incremental training strategy is adopted—only the newly added data in the window is used to update the model parameters, reducing training time by 70% compared to full retraining.

3.4 Hierarchical Contribution Degree Quantification Method

To identify key emission drivers at each level, the Local Interpretable Model-agnostic Explanations (LIME) method is used to quantify the contribution degree of each factor. For each level, LIME calculates the importance score of each input factor (e.g., frequency for chip level, PSU efficiency for server level) by perturbing the factor value and observing the change in emission estimation results. This method provides interpretable contribution rankings, guiding targeted carbon reduction strategies. In this study, the perturbation range of each factor is set based on industry standards: for example, the frequency perturbation range is ±20% of the nominal value, and the PSU efficiency perturbation range is ±5%. The importance score is normalized to 0–1, with a score greater than 0.5 indicating a key emission driver. This quantification method not only helps identify high-impact factors but also enables sensitivity analysis—for example, a 10% reduction in GPU frequency leads to a 16.3% reduction in server emissions, as shown in the subsequent experimental results.

4 Multi-Scene Experimental Validation and Analysis

4.1 Experimental Design and Data Sources

The experiment uses a 14-month dataset (January 2024 to February 2025) from a commercial cloud datacenter in Northern China, which hosts LLM inference, cloud storage, and online computing services. The dataset includes:

Chip-Level Data: 5-minute interval measurements of CPU/GPU frequency, voltage, and utilization for 150 heterogeneous servers (70 GPU servers, 80 CPU-only servers) using Intel RAPL and NVIDIA NVML.

Server/Rack-Level Data: Power consumption of servers, PDUs, and cooling systems measured via IPMI and smart power meters, covering 20 racks with different server densities.

Datacenter-Level Data: Campus-level PUE (average 1.22), grid carbon intensity (0.58 kg CO_2/kWh), and renewable energy ratio (20% wind power) from the datacenter's energy management system.

Three baseline methods are selected for comparison: (1) PUE-based macro model; (2) chip-level physics model; (3) random forest-based data-driven model. The evaluation metrics include estimation accuracy, MAE, and hierarchical contribution degree.Before model training, the dataset undergoes a strict preprocessing process: first, abnormal values (e.g., sensor data exceeding the normal range) are removed using the 3σ criterion; second, missing values are filled using linear interpolation, as missing data accounts for less than 2% of the total dataset; third, all features are normalized to the range of 0 – 1 to accelerate model convergence. The experiment is conducted on a server with two Intel Xeon Gold 6348 CPUs and four NVIDIA A100 GPUs, using Python 3.9 and deep learning frameworks such as PyTorch and DGL (Deep Graph Library) for model implementation.

4.2 Accuracy Comparison and Analysis

Table 1 shows the performance of HCCEEM and baseline models on the test set (20% of the dataset, 2.8 months of data):

Table 1. Performance of HCCEEM and baseline models on the test set

Model	Estimation Accuracy (%)	MAE (kg CO_2/kWh)
PUE-Based Macro Model	69.3	0.092
Chip-Level Physics Model	83.5	0.054
Random Forest-Based Model	87.9	0.038
HCCEEM (Proposed)	95.7	0.021

Key findings are as follows:

1.HCCEEM achieves the highest estimation accuracy (95.7%), outperforming the PUE-based model by 26.4% and the random forest model by 7.8%. This is due to the GNN-based cross-level aggregation, which captures both hardware heterogeneity and infrastructure energy losses.

2.The MAE of HCCEEM is only 0.021 kg CO_2/kWh, a 77.2% reduction compared to the PUE-based model, demonstrating its superior precision for fine-grained carbon accounting.

3.During peak LLM inference periods, HCCEEM maintains an accuracy of 94.1%, while the random forest model's accuracy drops to 81.2%, validating the effectiveness of the real-time adaptation mechanism.

To further verify the model's adaptability, additional experiments are conducted on three sub-datasets

corresponding to different workload scenarios: LLM inference, cloud storage, and online computing. The results show that HCCEEM achieves accuracy of 94.1%, 96.2%, and 95.5% in these three scenarios, respectively, while the random forest model's accuracy varies significantly (81.2%, 90.5%, 88.7%). This indicates that HCCEEM has strong stability across different workloads. In addition, a scalability test is performed by gradually increasing the number of servers in the dataset from 50 to 150; the results show that HCCEEM's accuracy decreases by only 1.2%, while the random forest model's accuracy decreases by 5.3%, proving that HCCEEM is suitable for large-scale datacenter clusters.

4.3 Hierarchical Contribution Degree Empirical Analysis

Using the LIME method, the contribution degree of each level to total emissions is quantified:

1.Chip Level: Contributes 65.2% of server-level emissions, with CPU/GPU utilization (importance score 0.72) and frequency (0.68) as the top two drivers. Reducing GPU frequency by 10% during low-load periods can reduce server emissions by 16.3%.

2.Server Level: Contributes 80.5% of rack-level emissions, with server type as the key factor—GPU servers contribute 4.1x more emissions than CPU-only servers. Optimizing server workload allocation (e.g., migrating non-critical tasks to CPU servers) can reduce rack emissions by 22.1%.

3.Rack Level: Contributes 90.2% of datacenter-level emissions, with cooling loss factor as the dominant driver (importance score 0.75). Adopting liquid cooling technology for high-density racks can reduce cooling losses by 40%, cutting datacenter emissions by 13.5%.

5 PRACTICAL IMPLICATIONS AND FUTURE DIRECTIONS

5.1 Practical Implications for Carbon Reduction

HCCEEM provides targeted optimization guidance for different stakeholders:

Hardware Engineers: Focus on chip-level dynamic power management, such as adjusting frequency and voltage based on task type, to reduce the core emission source of servers.

Rack Operators: Optimize server density and cooling strategies, such as adopting liquid cooling for high-density GPU racks, to minimize rack-level energy losses.

Datacenter Managers: Increase the proportion of renewable energy and optimize workload scheduling across racks, to reduce the amplification effect of energy structure on emissions.

For regulatory authorities, HCCEEM can serve as a standard tool for carbon report verification, enabling them to accurately assess the carbon emission data submitted by enterprises. For small and medium-sized datacenters, the lightweight version of HCCEEM (which simplifies the GNN module and reduces data collection requirements) can provide cost-effective carbon estimation solutions, helping them comply with carbon regulations without heavy investment. In addition, HCCEEM can be integrated into datacenter energy management systems (EMS) to realize real-time carbon monitoring and optimization, forming a closed-loop of "estimation-optimization-verification" for carbon reduction.

5.2 Limitations and Future Research Directions

The current framework has three limitations: first, it requires multi-level data collection, which is costly for small and medium-sized datacenters; second, it is validated only on cloud datacenters, and needs to be adapted for edge datacenters with distributed energy supply; third, it does not consider the carbon emissions from hardware manufacturing and disposal.

Future research will focus on three directions: (1) developing a lightweight version of HCCEEM for small datacenters, reducing data collection requirements; (2) adapting the framework to edge datacenters, integrating distributed renewable energy factors; (3) expanding the model to cover the entire lifecycle of datacenter hardware, from manufacturing to disposal. In addition, future research will explore the combination of HCCEEM with digital twin technology, constructing a virtual replica of the datacenter to simulate the impact of different carbon reduction measures (e.g., server workload scheduling, cooling system upgrades) on emissions. Another promising direction is to apply federated learning to HCCEEM, enabling multiple datacenters to train the model collaboratively without sharing sensitive data, which can address the data scarcity problem while protecting data privacy.

6 CONCLUSION

This study proposes a hierarchical coupling carbon emission estimation model for datacenters, which bridges the gap between macro-level rough calculation and micro-level isolated measurement. By constructing a four-level traceability chain, integrating physics-based modeling and GNN-based aggregation, and designing a real-time load adaptation mechanism, HCCEEM achieves high estimation accuracy and interpretability. Experimental results show that the model outperforms traditional methods in both stable and peak load scenarios, and provides fine-grained hierarchical contribution degree analysis to guide targeted carbon reduction. This research not only advances the theoretical

research of datacenter carbon accounting, but also provides a practical tool for the digital industry to achieve carbon neutrality goals. With the continuous growth of digital economy, HCCEEM is expected to be widely applied in cloud, edge, and hybrid datacenters, promoting the sustainable development of the global digital industry. In the long term, the hierarchical coupling idea proposed in this study can also be extended to other high-energy-consuming industries, such as manufacturing and transportation, providing a universal framework for accurate carbon emission estimation and reduction.

ACKNOWLEDGMENTS

Not Applicable.

FUNDING

Not Applicable.

INSTITUTIONAL REVIEW BOARD STATEMENT

Not Applicable.

INFORMED CONSENT STATEMENT

Not Applicable.

DATA AVAILABILITY STATEMENT

Not Applicable.

CONFLICT OF INTEREST

Not Applicable.

PUBLISHER'S NOTE

All claims expressed in this article are solely those of the authors and do not necessarily represent those of their affiliated organizations, or those of the publisher, the editors and the reviewers. Any product that may be evaluated in this article, or claim that may be made by its manufacturer, is not guaranteed or endorsed by the publisher.

AUTHOR CONTRIBUTIONS

Not application.

ABOUT THE AUTHORS

LIU, Wenwen

Bytedance, CN, liuwenwen.jessica@bytedance.com.

REFERENCES

[1] The Green Grid. (2024). Global Datacenter Energy Efficiency Report 2024. Beaverton, OR: The Green Grid Association.

[2] Zhang, H., Wang, Y., & Li, J. (2024). Dynamic PUE Correction for Datacenter Carbon Emission Estimation. IEEE Transactions on Green Communications and Networking, 8(3), 2145–2158.

[3] Wang, C., Li, S., & Zhao, H. (2023). Chip-Level Dynamic Power Consumption Modeling for GPU Servers. IEEE Transactions on Computer-Aided Design of Integrated Circuits and Systems, 42(8), 2678–2691.

[4] Chen, Y., Liu, X., & Zhang, Q. (2024). Random Forest-Based Carbon Emission Prediction for Cloud Datacenters. Sustainable Computing: Informatics and Systems, 42, 100923.

[5] Intel Corporation. (2023). Intel RAPL Technology Specification. Santa Clara, CA: Intel Press.

[6] NVIDIA Corporation. (2024). NVIDIA NVML API Guide. Santa Clara, CA: NVIDIA Developer Network.

[7] Li, Q., Yang, R., & Zhang, L. (2024). Edge Datacenter Carbon Emission Modeling with Distributed Renewable Energy Integration. IEEE Internet of Things Journal, 11(5), 8921–8934.

[8] Zhang, M., Chen, J., & Wang, T. (2023). Lifecycle Carbon Footprint Assessment for Datacenter Hardware: From Manufacturing to Disposal. Journal of Cleaner Production, 426, 138912.

[9] Wang, Y., Gao, S., & Li, X. (2024). Graph Neural Networks for Hierarchical Energy Efficiency Optimization in Heterogeneous Datacenters. IEEE Transactions on Parallel and Distributed Systems, 35(2), 567–579.

[10] Chen, G., Zhou, Y., & Xu, C. (2023). Liquid Cooling Technology for High-Density GPU Racks: Energy Saving and Carbon Reduction Potential. Applied Energy, 345, 121876.

[11] Liu, H., Sun, W., & Zhao, Y. (2024). Renewable Energy Scheduling for Cloud Datacenters: A Carbon-Aware Optimization Approach. Energy Conversion and Management, 312, 117345.

[12] Zhao, J., Huang, L., & Zhu, M. (2023). Dynamic Workload Scheduling for Carbon Reduction in Multi-Tenant Datacenters. IEEE Transactions on Cloud Computing, 11(4), 3210–3223.

[13] Huang, Z., Wu, X., & Liu, S. (2024). Lightweight Carbon Emission Estimation Model for Small-Scale Datacenters with Limited Monitoring Data. IEEE Access, 12, 45678–45692.

[14] International Energy Agency (IEA). (2024).

Digitalization and Energy Report 2024. Paris: International Energy Agency.