Engineering xxx (xxxx) xxx

Contents lists available at ScienceDirect

Engineering

journal homepage: www.elsevier.com/locate/eng

Research
Engineering Management—Perspective

Building Accurate Energy-Use Statistics for Data Centers

Yong-Zhen Wang [a,b], Te Han [b,c,d,e,*], Yi-Ming Wei [b,c,d,e,*]

[a] School of Mechanical Engineering, Beijing Institute of Technology, Beijing 100081, China
[b] Beijing Lab for System Engineering of Carbon Neutrality, Beijing Municipal Education Commission, Beijing 100081, China
[c] Center for Energy and Environmental Policy Research, Beijing Institute of Technology, Beijing 100081, China
[d] School of Management, Beijing Institute of Technology, Beijing 100081, China
[e] NSFC Basic Science Center for Energy and Climate Change, Beijing 100081, China

ARTICLE INFO

Article history:
Received 30 April 2025
Revised 27 November 2025
Accepted 7 December 2025
Available online xxxx

ABSTRACT

With the rapid expansion of cloud computing and large-scale artificial intelligence models, building accurate and transparent energy-use statistics for data centers has become a critical challenge for global energy systems and climate governance. Existing studies report strikingly divergent estimates of global data center electricity consumption, ranging from 196 to 1200 TW·h in 2020, a more than sixfold difference. Such discrepancies reveal profound uncertainties and structural deficiencies in current energy accounting frameworks. Conventional estimation approaches rely heavily on indirect assumptions, proxy indicators, or highly aggregated regional and national statistics, obscuring the true electricity demand of data centers. This lack of statistical transparency distorts energy and carbon accounting, weakens power system planning, and constrains the effective integration of renewable energy with rapidly growing computing demand. This paper highlights that data centers should be treated as a distinct and strategically important end-use energy sector. It emphasizes the need for grid-informed energy registration, enhanced artificial intelligence identification techniques to improve the accuracy and verifiability of energy statistics. Furthermore, the paper emphasizes that policymakers should establish coordinated policy frameworks, enforce standardized energy reporting, and design appropriate incentive mechanisms to encourage data centers to participate in demand response programs and electricity markets, thereby unlocking load flexibility and supporting a secure, low-carbon energy transition.

1. Introduction

Data centers have emerged as the backbone of the modern digital economy, enabling a wide range of services, from cloud computing to advanced artificial intelligence (AI) applications. Over the past decades, global computing capacity has grown significantly, driving a corresponding rise in the energy consumption of data centers. In particular, the rapid rise of large AI systems has further heightened these concerns [1]. The computing power required for training and running these models contributes substantially to the growing energy demands of data centers. According to a recent report, the energy demands of large AI systems and data centers could soon rival the energy consumption of entire nations [2].

As data centers expand in both scale and influence, their growing energy consumption increasingly underscores the need for accurate energy accounting and resource measurement [3]. Reliable data is essential not only for achieving low-carbon and sustainable development but also for optimizing industry planning and more effectively aligning renewable energy generation with computing demands. Recent findings published in *Joule* demonstrate substantial uncertainties in global estimates of data center energy consumption. For example, in 2020, estimates ranged from as low as 196 TW·h to as high as 1200 TW·h, reflecting a difference of more than sixfold [4]. These discrepancies present substantial obstacles to informed strategic planning as well as the sustainable development of data centers and their affiliated digital industries. Recent estimates from the European Commission suggest that European Union (EU) data centers may consume between 98.5 and 160 TW·h by 2030, contingent upon modeling assumptions and definitions [5]. This broad range indicates that regional energy forecasts remain highly sensitive to methodological variation.

2. Challenges in estimating data center energy consumption

Accurately estimating the energy consumption of data centers remains a significant ongoing challenge for researchers, grid operators, and policymakers alike. As highlighted by recent studies in

---

[*] Corresponding authors.
E-mail addresses: hante@bit.edu.cn (T. Han), wei@bit.edu.cn (Y.-M. Wei).

https://doi.org/10.1016/j.eng.2025.12.014

Please cite this article as: Y.-Z. Wang, T. Han and Y.-M. Wei, Building Accurate Energy-Use Statistics for Data Centers, Engineering, https://doi.org/10.1016/j.eng.2025.12.014

Science and *Joule* [1,4], mainstream estimation methods can be broadly categorized into two principal paradigms: bottom-up approaches, which infer total energy use from infrastructure-level parameters such as server utilization, rack density, and power usage effectiveness (PUE) [6]; and top-down approaches, which derive aggregate electricity consumption from national statistics, industry surveys, or utility-report billing data.

2.1. Bottom-up estimation: parameter-based but unreliable

Bottom-up modeling estimates overall energy consumption by integrating equipment-level performance metrics with data center infrastructure parameters such as rack density, server utilization, PUE, and the total number of racks [4,6]. Although conceptually straightforward, this parameter-based method is characterized by extensive uncertainties and limited access to accurate operational datasets, which collectively undermine its reliability for comprehensive system-level energy accounting [7].

Fundamental metrics employed in bottom-up estimation vary considerably across facilities. Rack density, defined as the ratio of occupied cabinets to total designed capacity, fluctuates depending on data center scale, client requirements, and reserved capacity for future expansion. Utilization rates of information technology (IT) equipment are highly dynamic, ranging from less than 10% to nearly 100% depending on workload intensity [8]. PUE further accentuates this variability, with values spanning approximately 1.1 in hyperscale facilities with advanced cooling systems to over 2.0 in small or medium-sized centers using room-level air conditioning [9,10]. Such fluctuations render it extremely challenging to establish consistent or comparable energy models across sites. Beyond these intrinsic variations, resource allocation practices also introduce additional uncertainty. Many data centers implement resource-sharing and capacity-reservation schemes to accommodate diverse client demands and future expansion, resulting in persistent mismatches between available capacity and actual utilization. The uncertainties associated with parameters within the bottom-up estimation framework can be amplified across multiple layers of calculation, significantly undermining the reliability of final energy consumption estimates.

Some studies have attempted to estimate data center energy consumption using server shipment statistics and computing energy efficiency benchmarks such as SPECpower_ssj2008 [6,11]. However, these proxy metrics frequently fail to reflect the inherent heterogeneity of real-world operations. Server shipments do not accurately represent actual utilization, as many units remain idle, underused, or retired, resulting in pronounced deviations between theoretical and operational energy use. The SPECpower_ssj2008 benchmark, established by the Standard Performance Evaluation Corporation (SPEC), is a widely recognized standardized protocol for evaluating the performance-per-watt efficiency of computer servers. It measures the number of Java operations (SSJ_ops) executed per watt of power consumed, providing a normalized and comparable indicator of computing energy efficiency. As illustrated in Fig. 1, typical Intel and AMD servers have achieved nearly a tenfold improvement in SSJ_ops per watt between 2008 and 2024, reflecting remarkable hardware progress. Nevertheless, these benchmark results are obtained under highly controlled laboratory conditions and cannot capture the diversity of real-world workloads, heterogeneous hardware configurations, or the coexistence of outdated equipment. Consequently, depending on such static metrics as proxies for operational performance introduces systematic bias, compromises estimation accuracy, and limits the reliability of bottom-up energy estimation, particularly when applied across heterogeneous facilities.

Finally, bottom-up estimation relies heavily on self-reported or internally aggregated enterprise data, which often suffer from limited transparency and consistency. Under cost or increasing regulatory pressure, some operators may underreport consumption, resulting in significant discrepancies among studies. Without independent verification, these datasets are inadequate for robust energy governance.

In summary, while bottom-up estimation offers intuitive and equipment-oriented insights, its reliance on uncertain parameters and opaque reporting mechanisms produces fragmented and unreliable assessments [1]. These shortcomings underscore the need for complementary, system-level monitoring mechanisms that can provide objective and standardized energy accounting across diverse data center types.

2.2. Top-down estimation: macro-level but underdeveloped

Top-down modeling estimates total energy consumption using aggregated regional or national data, typically sourced from government statistics, industry reports, or large-scale surveys [4,6]. The primary advantage of this approach lies in its reliance on directly measured values, such as total electricity metering data and recorded network traffic volumes, enabling a more accurate system-level assessment of energy use. Compared with bottom-up approaches, it offers broader coverage and enhanced objectivity, since it does not depend on self-reported enterprise data. Consequently, it has become the standard methodology for evaluating sectoral energy demand at regional or national scales.

However, this macro-level approach exhibits two key limitations in estimating data center energy use. First, the underlying statistical data is often incomplete. Aggregated electricity records cannot precisely distinguish the portion of power consumption attributable to data centers in mixed-use or embedded facilities where meters are shared with other operations. Second, there is a notable absence of standardized registration and classification frameworks. In official energy statistics, data centers are frequently categorized under broad information and communication technology (ICT) or commercial electricity categories. These limitations introduce systematic biases and impede the reliable quantification of data center energy consumption across national or regional estimates.

To enhance accuracy from the grid perspective, a targeted monitoring strategy is required: AI-based identification can detect existing or embedded facilities that remain unregistered, while registration and reporting protocols should be strengthened for newly constructed data centers.

3. Governing AI via AI

In existing data centers, a persistent challenge for top-down energy accounting is the absence of explicit energy-use labels, particularly in mixed-use environments such as commercial complexes or campuses. While large-scale facilities typically implement advanced submetering and PUE monitoring systems, small and medium-sized data centers—particularly those embedded within multi-purpose buildings—frequently lack such infrastructure. Consequently, grid operators cannot reliably isolate their electricity consumption. With the increasing deployment of smaller edge data centers, driven by lightweight AI workloads such as DeepSeek, this fragmentation further complicates system-level energy statistics. In cases where submetering is unavailable, emerging AI-based approaches (Fig. 2) provide a scalable alternative by identifying data-center-specific load patterns from aggregated signals.

AI-driven classification models can analyze historical power data and facility attributes to infer the presence and intensity of data-center operations. Studies have shown that power-based

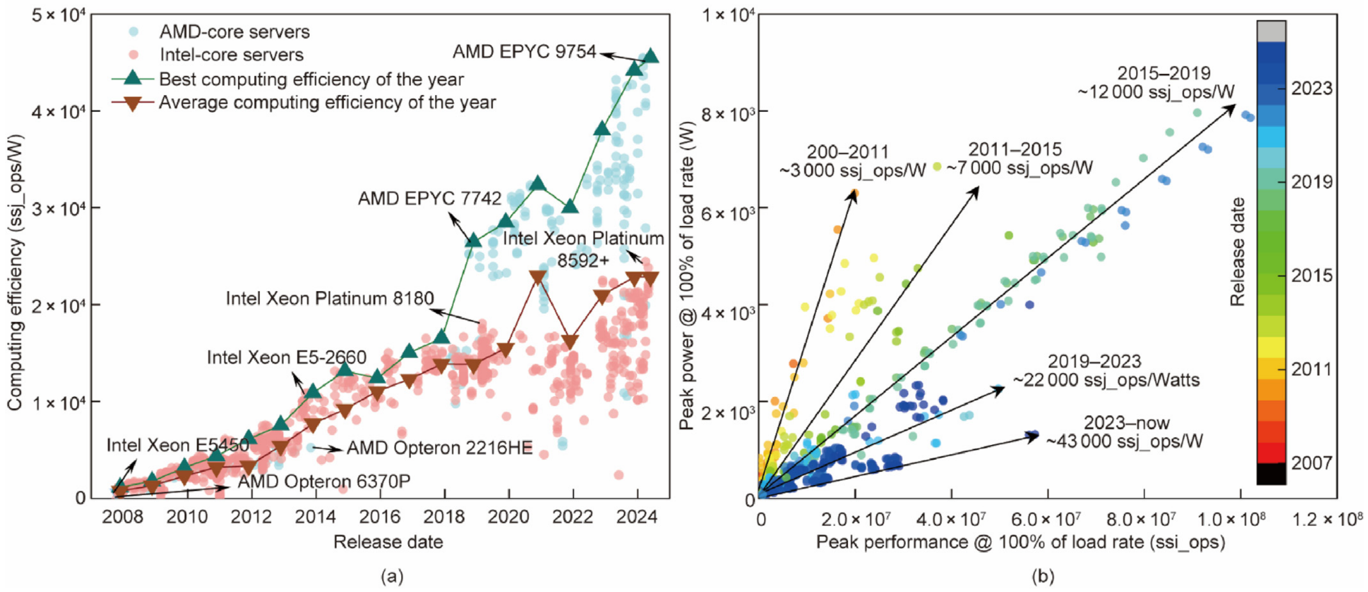

Fig. 1. Benchmarking results from SPECpower_ssj2008 for server computing energy efficiency. (a) Evolution of computing energy efficiency for typical servers from 2007 to 2024, and (b) corresponding changes in server power consumption and performance from 2007 to 2024.

classifiers and hybrid deep-learning models can achieve accuracy levels exceeding 85%–90% in distinguishing data-center workloads from other building types [12,13]. As presented in Fig. 2, the blue section at the top represents several buildings, each equipped with an independent electricity meter; however, the specific functions of the buildings remain unidentified. To address this limitation, AI-driven classification models analyze aggregated power consumption signals, enabling the extraction of load profiles that correspond specifically to data center operations. These models leverage the operational characteristics of data centers—continuous 24/7 operation, high load density, and relatively stable time-series profiles—which contrast sharply with the highly variable, occupancy-driven patterns of offices, malls, or residential complexes. Even in AI-intensive or High-Performance Computing (HPC) facilities, where workloads may fluctuate dynamically, the overall signal remains structured and predictable, providing a robust foundation for AI-based inference.

When direct metering is unavailable, these classification principles can be extended using non-intrusive load monitoring (NILM). The green section at the bottom of Fig. 2 presents a scenario in which multiple buildings share a single electricity meter, producing aggregated or mixed power consumption data. In such cases, NILM techniques disaggregate device- or facility-level loads from aggregated consumption by detecting characteristic load events (e.g., server startup, cooling cycles) using time-series models such as long short-term memory (LSTM) or Transformer architectures.

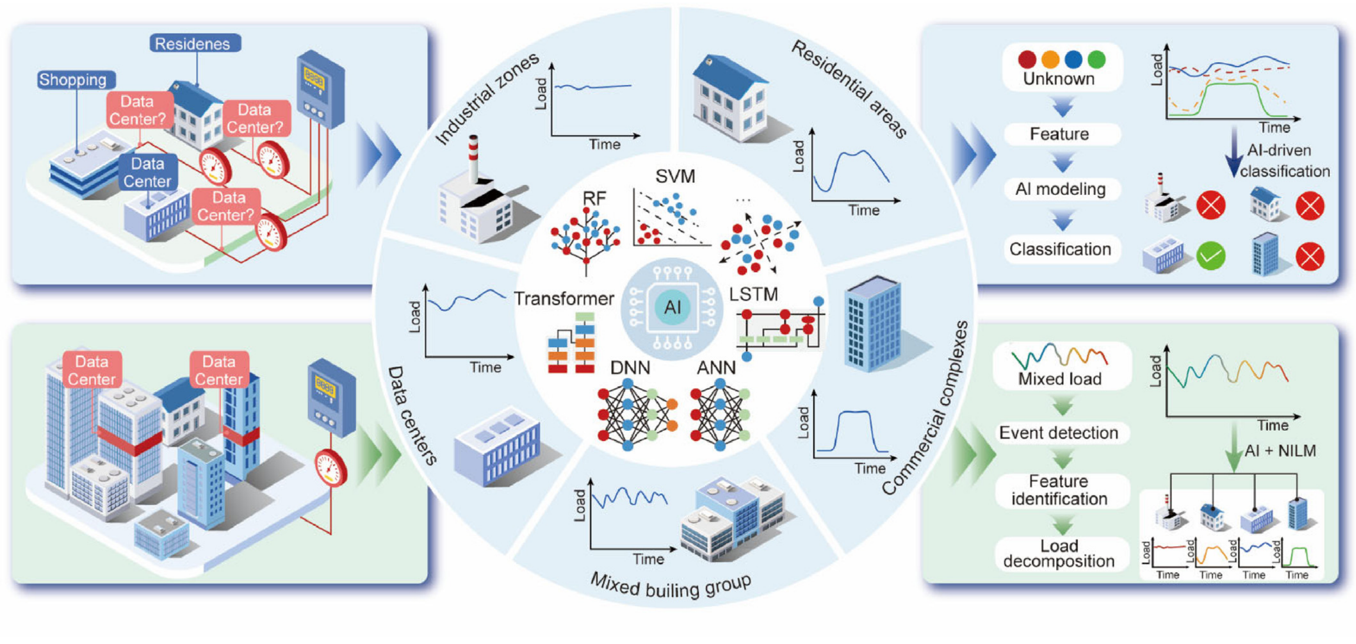

Fig. 2. Leveraging advanced AI methodologies to identify consumption patterns specific to data centers and dissect aggregated load data. LSTM: long short-term memory; DNN: deep neural networks; ANN: artificial neural networks; RF: random forest; SVM: support vector machine. NILM: non-intrusive load monitoring.

Recent studies have demonstrated high device-level load identification accuracy in commercial and industrial contexts [14,15], confirming the feasibility of applying these techniques to identify data center loads in mixed-use environments. In essence, AI-based NILM provides a non-invasive framework to reconstruct fine-grained energy attribution in situations where physical metering is impractical.

To improve the robustness of energy estimation in real-world deployments, regions with high energy consumption per unit area and relatively stable load profiles should be prioritized, as these characteristics are indicative of typical data center operations. Furthermore, dynamic validation through small-scale manual labeling or audit checks can improve model performance, ensuring AI systems remain adaptive and accurate even when monitoring data is low resolution or noisy. This integrated methodology, incorporating AI with NILM and human-in-the-loop feedback, provides a practical and scalable framework for implementing energy governance in existing infrastructures, especially in mixed-use buildings where direct metering is limited.

4. Reforming data center energy registration

When establishing new data center infrastructure, a critical requirement is the formal registration of energy usage with the grid. While AI technologies can play a key role in estimating energy consumption for existing data centers, enhancing the registration process for new data centers remains a fundamental step. First, the data collected during registration serves as a critical resource for optimizing and validating AI models. This information provides valuable prior knowledge. In comparison to relying solely on AI-driven identification, detailed registration data, including building purpose, rack configuration, and energy usage patterns, substantially reduces uncertainty and improves the accuracy of energy consumption estimations.

Moreover, registration data plays a pivotal role in policy implementation and industry planning. Clear energy-use labeling enhances regulatory transparency, resolves ambiguities between data centers and other building types, and mitigates statistical blind spots. Analogous to other major terminal energy-consuming sectors (e.g., steel, cement, glass), the establishment of a comprehensive registration mechanism will ultimately enable governments to publicly disclose data center energy consumption and carbon emissions, thereby strengthening sectoral oversight and informing strategic planning.

4.1. Centralized vs decentralized registration

In regions with centralized grid operations, such as China, the unified management structure of major operators (e.g., State Grid and Southern Grid) provides favorable institutional conditions for standardized energy registration and reporting. However, despite this advantage, few top-down registration mechanisms have been successfully implemented in China. Data center development in China is overseen by multiple agencies, including the National Development and Reform Commission, the Ministry of Industry and Information Technology of the People's Republic of China, and regional Communications Administrations, whereas grid operators are often excluded from early-stage planning. Consequently, critical energy-related information is seldom shared with grid companies in a timely or standardized manner. As project approvals frequently precede grid infrastructure planning, operators typically engage only at the load connection stage, depending on self-declared power demand rather than detailed technical information such as rack counts or cooling configurations. This institutional fragmentation constrains the grid's capacity to anticipate, monitor, and manage the expanding energy footprint of data centers from the outset. To address this gap, enhancing data disclosure during project registration, including key indicators such as building area, rack quantity, and equipment type, could substantially improve the transparency and accuracy of national data center statistics.

In regions with decentralized grid operations, such as the United States and Europe, the presence of multiple independent grid operators often results in fragmented oversight and inconsistent implementation. While large, dedicated data centers are generally registered and recognized by utilities, identifying smaller or embedded facilities—particularly those within mixed-use buildings—remains challenging. Regulatory frameworks also vary widely across member states. For example, the Netherlands enforces stringent zoning and mandatory requirements for data centers [16], whereas many other EU countries impose no comparable obligations, resulting in uneven data quality across the region. In this context, combining market incentives with technical support offers a more pragmatic approach to improved registration. Financial mechanisms, including green electricity pricing, renewable energy prioritization, or certification schemes, can encourage operators to voluntarily disclose detailed energy information. Such incentive-aligned approaches enhance corporate competitiveness while simultaneously improving data completeness and transparency.

Overall, the organization of grid operations fundamentally shapes the design of registration mechanisms. Centralized systems benefit from institutional integration and procedural standardization, whereas decentralized systems may rely more often on incentive-based and interoperable data governance frameworks. Recognizing these structural differences is essential for developing scalable and adaptive registration models that align with regional electricity architectures.

4.2. Integrating reporting frameworks into grid-informed registration

The European experience provides valuable insights for improving data center energy registration. The European Code of Conduct for Energy Efficiency in Data Centers (2008) pioneered the voluntary disclosure of key indicators, including total energy consumption, PUE, and rack utilization [17]. Building on this foundation, the European Commission introduced a binding sustainability reporting framework in 2024 [18], requiring large facilities (currently 500 kW and above [19]) to disclose standardized metrics such as total energy consumption, renewable energy share, and cooling efficiency [20,21]. These initiatives constitute a notable step toward enhanced transparency and comparability across operators. However, these frameworks remain largely statistical and ex post, reporting aggregated energy data periodically, often using coarse categories without consistent identifiers or direct linkage to electricity meters. While such designs can reveal sectoral trends, they are insufficient for fine-grained, auditable accounting, particularly for small or embedded data centers sharing infrastructure with other facilities.

To improve statistical accuracy, registration should evolve from disclosure to grid-informed registration. Specifically, each facility record should include: ① a standard technical descriptor set (e.g., rated IT power, cooling type, redundancy, rack count, floor area); ② a uniform building/use taxonomy (dedicated vs embedded/mixed-use); and ③ a direct link to the facility's metered time series within grid information systems. This design allows consistent facility-level attribution, transforming static reports into auditable, high-resolution energy statistics.

The optimization of energy registration can be integrated within the existing grid management framework, without imposing substantial additional burdens on operators. By incorporating energy classification and registration requirements into the existing procedures for grid connection and electricity application, data centers can be clearly distinguished from conventional buildings at the point of connection. When combined with the widespread deployment of smart meters, this approach facilitates more accurate and consistent data collection.

5. Energy estimation reveals flexibility, flexibility brings benefits

Integrating top-down energy registration and consumption statistics for data centers inevitably imposes additional regulatory burdens and costs on grid operators. Nevertheless, this approach enables more accurate and verifiable energy statistics. By linking standardized registration with metered data directly to the grid, operators can identify and account for flexible power resources associated with data centers. In doing so, the grid not only gains a clearer understanding of data center energy usage but also unlocks potential synergies between computing power and electricity supply [22].

From the grid operator's perspective, data centers resemble traditional industrial loads, such as steel and cement plants, especially in terms of high energy consumption. Furthermore, data centers, particularly those focused on HPC and AI tasks, demonstrate substantial flexibility temporally and spatially [23,24]. As illustrated in Fig. 3, the green and blue task transfer arrows and blocks indicate that data center power loads fluctuate in real time based on computing tasks. Batch processing operations, such as large-scale model pre-training that typically allow delays of several hours, provide scheduling flexibility for the grid. Studies suggest that nearly 40% of batch processing tasks can be dynamically adjusted to reduce operational costs [25,26]. This inherent flexibility positions data centers as critical assets for grid-level energy management, with considerable potential to integrate with renewable energy infrastructures and support grid stability. Additionally, tasks can be transferred between geographically distributed data centers through fiber-optic networks, allowing instantaneous power transfer across locations and creating opportunities to align computing workloads with available clean energy sources.

Therefore, these adaptable resources can also be integrated into the monitoring and management framework. Mechanisms such as sustainable electricity pricing and demand-response programs can incentivize data centers to actively engage in grid load optimization. In this manner, the dual approach—tracking energy consumption while regulating flexibility—can deliver substantial benefits to the electricity network. Quantitative evidence from recent studies underscores the economic advantages of leveraging data center flexibility. For example, participation in demand response programs by data centers has been shown to reduce energy purchase costs by up to 24.19% [27]. Model-based analyses indicate that a 1% increase in temporally and spatially flexible loads from data centers can lower costs associated with renewable energy storage, scheduling, and optimization by approximately (1.29 ± 0.07) $\mathrm{EUR}\cdot(\mathrm{MW}\cdot\mathrm{h})^{-1}$ [28].

Strengthening top-down energy registration is not only crucial for establishing accurate energy statistics but also fundamental for enabling the scalable detection and management of flexibility resources. By providing standardized, system-wide visibility into energy use and infrastructure characteristics, such a framework can facilitate demand-side flexibility across heterogeneous data center environments. In doing so, it establishes the institutional basis for data centers to evolve from passive electricity consumers into proactive contributors to grid stability and renewable energy integration.

6. Conclusions and recommendations

Owing to its growing scale, data center energy consumption requires greater scrutiny and should be considered as a distinct

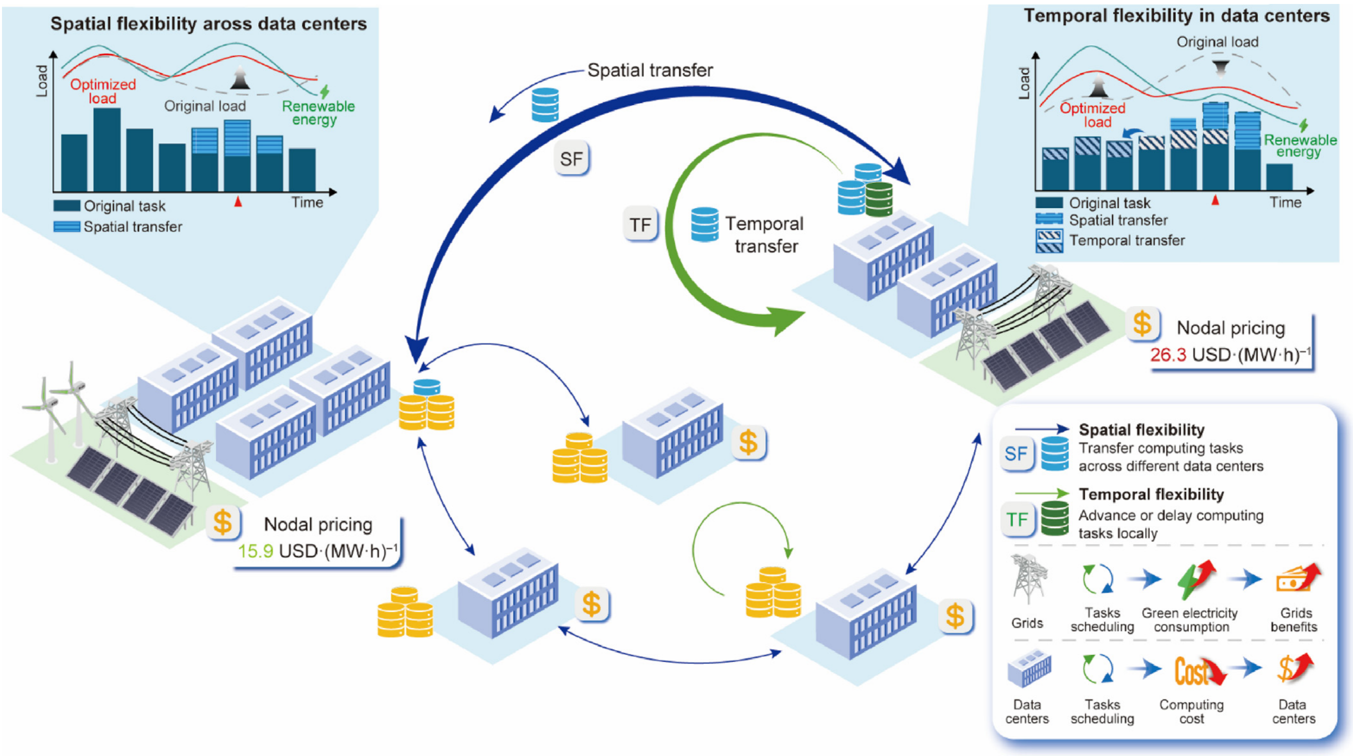

Fig. 3. Temporal and spatial flexibilities of computing power loads provide advantages for both grids and data centers. The blue bidirectional arrows denote spatial transfer of tasks between different data centers, while the green arrows represent temporal scheduling or shifting of tasks within a local data center.

category in energy management. Traditional bottom-up energy estimation methods for data centers encounter considerable challenges, primarily owing to the dynamic variations in rack density, utilization rates, and other factors, as well as limited data transparency, resulting in significant uncertainty in energy calculation and reporting. From the grid's perspective, smart meters can directly record electricity consumption data from data centers, providing a robust foundation for macro-level energy statistics. For existing unregistered data centers, it is critical to leverage big data analytics to develop intelligent identification algorithms based on load curve characteristics. By analyzing sustained high-load patterns and distinctive variability behaviors typical of data centers, these algorithms can accurately detect and distinguish data center energy consumption profiles without requiring direct metering, thereby allowing intelligent recognition and estimation of energy use across existing facilities. For newly established data centers, optimizing the energy registration process is crucial. This includes refining energy categorization, enhancing registration transparency, and implementing traceable reporting protocols to ensure that energy consumption data is clearly recorded and accessible for regulatory oversight and planning purposes.

Although implementing top-down energy estimation methods for data centers entails additional investments in infrastructure and operational costs, the potential economic advantages and system optimization derived from accurate data collection are substantial. By precisely monitoring energy consumption, the grid can identify and leverage flexible power resources from data centers, such as computing workload redistribution, which align more effectively with the electricity supply. This integration between computing power and electricity availability enables more efficient energy utilization across a broader scale, benefiting both data centers and the grid. For example, grid operators can use demand response mechanisms to adjust data center energy loads during peak demand periods, or increase computing the loads during periods of abundant renewable energy, thus optimizing wind and solar power utilization. This coordination also unlocks new revenue opportunities for grid operators—such as peak-shaving services and market transactions—whereas data centers can lower electricity costs through flexible energy management, creating mutual benefits for both sectors.

To achieve these objectives, policymakers should foster inter-departmental collaboration to establish harmonized energy reporting standards and allocate targeted financial support for AI technologies and grid infrastructure upgrades. Long-term regulatory and incentive mechanisms will be essential to encourage data centers to disclose energy consumption data and actively participate in electricity spot markets. For instance, the government can offer tax incentives or subsidies to support the deployment of advanced energy monitoring and management technologies in data centers, while grid companies should strengthen collaboration with operators to jointly develop and operate flexibility resource management platforms, ensuring the effective implementation and continuous improvement of energy estimation methods.

Despite their considerable energy footprint, data centers deliver far more than just computing power. By enabling advanced ICT and AI applications, data centers help energy, manufacturing, agriculture, and construction adopt more sustainable, low-carbon operations through optimized resource use and improved productivity. Recognizing the dual role of data centers in both energy consumption and broader sustainability efforts, it is essential to reassess their broader potential as catalysts for environmental and economic benefits.

CRediT authorship contribution statement

Yong-Zhen Wang: Writing – review & editing, Methodology, Investigation, Formal analysis, Conceptualization. **Te Han:** Writing – original draft, Methodology, Investigation, Conceptualization. **Yi-Ming Wei:** Writing – review & editing, Validation, Supervision, Investigation, Funding acquisition.

Declaration of competing interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

Acknowledgments

This research was supported by the Beijing Natural Science Foundation (L241083) and the National Natural Science Foundation of China (72595833, 72595830, 72293605, and 52006114).

References

[1] Masanet E, Shehabi A, Lei N, Smith S, Koomey J. Recalibrating global data center energy-use estimates. Science 2020;367(6481):984–6.
[2] Crawford K. Generative AI's environmental costs are soaring—and mostly secret. Nature 2024;626(8000):693.
[3] Debus C, Piraud M, Streit A, Theis F, Götz M. Reporting electricity consumption is essential for sustainable AI. Nat Mach Intell 2023;5(11):1176–8.
[4] Mytton D, Ashtine M. Sources of data center energy estimates: a comprehensive review. Joule 2022;6(9):2032–56.
[5] Kamiya G, Bertoldi P. Energy consumption in data centres and broadband communication networks in the EU. Report. Luxembourg: Publications Office of the European Union; 2024.
[6] Masanet E, Lei N, Koomey J. To better understand AI's growing energy use, analysts need a data revolution. Joule 2024;8(9):2427–36.
[7] Tang G, Jiang W, Xu Z, Liu F, Wu K. NIPD: non-intrusive power disaggregation in legacy datacenters. IEEE Trans Comput 2017;66(2):312–25.
[8] De Vries A. The growing energy footprint of artificial intelligence. Joule 2023;7(10):2191–4.
[9] Lei N, Masanet E. Climate- and technology-specific PUE and WUE estimations for U.S. data centers using a hybrid statistical and thermodynamics-based approach. Resour Conserv Recycl 2022;182:106323.
[10] Pei Q, Yuan Y, Hu H, Wang L, Zhang D, Yan B, et al. Working smarter not harder: hybrid cooling for deep learning in edge datacenters. IEEE Trans Sustain Comput 2025;10(4):804–19.
[11] Standard Performance Evaluation Corporation (SPEC). SPECpower ssj® 2008 Results – Query [Internet]. Gainesville: Standard Performance Evaluation Corporation; undated [cited 2025 Nov 28]. Available from: https://www.spec.org/cgi-bin/osgresults?conf=power_ssj2008.
[12] Li Y, Hu H, Wen Y, Zhang J. Power data classification: a hybrid of a novel local time warping and LSTM. arXiv:1608.04171.
[13] Karimi M, Arab F. Machine learning based framework for estimation of data center power using acoustic side channel. arXiv:2008.02481.
[14] Tanoni G, Principi E, Squartini S. Non-intrusive load monitoring in industrial settings: a systematic review. Renew Sustain Energy Rev 2024;202:114703.
[15] Li Y, Yao R, Qin D, Wang Y. Lightweight federated learning for on-device non-intrusive load monitoring. IEEE Trans Smart Grid 2025;16(2):1950–61.
[16] Frijling R. EED audit and energy saving obligations for Dutch data centers [Internet]. Lincoln: EkkoSense; 2024 May 16 [cited 2025 Apr 3]. Available from: https://www.ekkosense.com/resources/industry-insight/eed-audit-and-energy-saving-obligations-for-dutch-data-centers.
[17] Avgerinou M, Bertoldi P, Castellazzi L. Trends in data centre energy consumption under the European code of conduct for data centre energy efficiency. Energies 2017;10(10):1470.
[18] European Commission. Commission adopts EU-wide scheme for rating sustainability of data centres [Internet]. Brussel: European Commission; 2024 Mar 15 [cited 2025 Apr 3]. Available from: https://energy.ec.europa.eu/news/commission-adopts-eu-wide-scheme-rating-sustainability-data-centres-2024-03-15_en.
[19] Energy Authority. Reporting from data centres to the European database has started. Jacksonville: Energy Authority; 2024 Oct 14 [cited 2025 Apr 3]. Available from: https://energiavirasto.fi/en/-/reporting-from-data-centres-to-the-european-database-has-started.
[20] European Commission. Commission Delegated Regulation (EU) 2024/1364 of 14 March 2024 on the first phase of the establishment of a common Union rating scheme for data centres [Internet]. Brussel: European Commission; 2024 Mar 15 [cited 2025 Apr 3]. Available from: https://eur-lex.europa.eu/eli/reg_del/2024/1364/oj.
[21] Zhu X, Jiang W, Liu F, Zhang Q, Pan L, Chen Q, et al. Heat to power: thermal energy harvesting and recycling for warm water-cooled datacenters. In: Proceedings of the 2020 ACM/IEEE 47th Annual International Symposium on Computer Architecture (ISCA); 2020 May 30–Jun 3; held virtually. New Yor City: IEEE Press; 2020. p. 405–18.

[22] Wang Y, Han Y, Shen J, Han K, Li J, Ji W, et al. Data center integrated energy system for sustainability: generalization, approaches, methods, techniques, and future perspectives. Innov Energy 2024;1(1):100014.

[23] Pei Q, Wang L, Zhang D, Yan B, Yu C, Liu F. InferCool: enhancing AI inference cooling through transparent, non-intrusive task reassignment. In: Proceedings of the 2024 ACM Symposium on Cloud Computing, SoCC'24; 2024 Nov 20–22; Redmond, WA, USA. New York City: Association for Computing Machinery (ACM); 2024. p. 487–504.

[24] Liu F, Zhou Z, Jin H, Li B, Li B, Jiang H. On arbitrating the power-performance tradeoff in SaaS Clouds. IEEE Trans Parallel Distrib Syst 2024;25(10):2648–58.

[25] Cao Y, Cheng M, Zhang S, Mao H, Wang P, Li C, et al. Data-driven flexibility assessment for internet data center towards periodic batch workloads. Appl Energy 2022;324:119665.

[26] Zeng B, Zhou Y, Xu X, Cai D. Bi-level planning approach for incorporating the demand-side flexibility of cloud data centers under electricity-carbon markets. Appl Energy 2024;357:122406.

[27] Han O, Ding T, Mu C, Jia W, Ma Z. Waste heat reutilization and integrated demand response for decentralized optimization of data centers. Energy 2023;264:126101.

[28] Riepin I, Brown T, Zavala VM. Spatio–temporal load shifting for truly clean computing. Adv Appl Energy 2025;17:100202.