International Journal of Science and Research Archive
eISSN: 2582-8185
Cross Ref DOI: 10.30574/ijsra
Journal homepage: https://ijsra.net/

(Review Article)

Check for updates

A stochastic optimization framework for AI-driven commissioning processes in data centers: Enhancing lifecycle efficiency and cost reduction

Nnadozie Odinaka [1, *], Martin Dillum [2] and Oghnetega Deborah Wash-Anigboro [3]

[1] *Scheller College of Business, Georgia Institute of Technology, USA.*
[2] *Marshall School of Business, University of Southern California, USA.*
[3] *Harbert College of Business, Auburn University, USA.*

International Journal of Science and Research Archive, 2025, 16(01), 567-574

Publication history: Received on 18 May 2025; revised on 05 July 2025; accepted on 08 July 2025

Article DOI: https://doi.org/10.30574/ijsra.2025.16.1.2048

Abstract

This study introduces a new AI-driven framework for commissioning hyperscale data centers, replacing traditional checklist methods with a more efficient, autonomous process. By integrating Bayesian optimization with a real-time digital twin, the system dynamically plans and adjusts performance tests, aiming to maximize efficiency while minimizing cost and risk. The approach uses Gaussian-process models to update its understanding as data is collected, enabling smarter decisions with less testing. Results show significant benefits over conventional methods, including 15–25% faster commissioning, lower upfront costs, and 8–12% energy savings over the data center facility's lifetime. The proposed framework therefore offers a scalable, AI-driven pathway to accelerate deployment, cut costs, and embed continuous optimization capabilities from day one in modern data-center infrastructure.

Keywords: Data-Center Commissioning; Bayesian Optimization; Stochastic Optimization; Digital Twin; Hyperscale Facilities; Lifecycle Efficiency; Cost Reduction; Artificial Intelligence

1. Introduction

Hyperscale data centers, which support cloud and AI services, require thorough commissioning to ensure reliability and efficiency from the start. Commissioning verifies that infrastructure systems (like cooling, power, and controls) function according to design and owner requirements as stated by Chakrabarty et al. (2021). When done properly, it enhances performance, reduces early equipment failures, and prevents costly outages (Dataknox, 2025). Additionally, it delivers financial and energy benefits, with commissioned buildings achieving a median energy savings of ~16% and a payback period of about 1.1 years (Synergy Engineers, 2022). Overall, commissioning is a vital process that mitigates risk and adds value to data center deployment.

According to Chen et al (2021), traditional commissioning is often seen as expensive and time-consuming, especially in hyperscale data centers that must scale quickly. These large facilities, with thousands of components, face significant risks from inefficient or delayed commissioning, which can cause major project delays—43% of which are linked to poor commissioning practices (Highjoule, 2023). Incomplete commissioning can also lead to costly failures, such as a $17 million outage in Tokyo caused by an unchecked logic error (Highjoule, 2023). These issues highlight the need for smarter, faster, and more cost-effective commissioning methods, including the use of advanced AI, to ensure reliability at scale.

* Corresponding author: Nnadozie Odinaka

This paper introduces a new stochastic optimization framework using Bayesian optimization to improve AI-driven commissioning in hyperscale data centers, specifically during pre-functional and functional performance testing. The goal is to boost efficiency and lower costs by optimizing and automating test planning under uncertainty. The paper reviews current commissioning practices and AI applications, highlights the underuse of Bayesian methods, presents a mathematical model and system architecture for AI-based commissioning, and outlines the expected benefits in cost and performance throughout the data center lifecycle.

1.1. Data Center Commissioning in Hyperscale Environments

The commissioning process in data centers is a multi-stage quality assurance procedure designed to validate the installation and performance of critical systems before and during operation (Mark H, 2025). It typically follows standardized protocols (e.g. ASHRAE Guidelines or Uptime Institute Tier certification processes) and includes five key stages as stated by Brown et al. (2015): factory acceptance testing, site inspection, pre-functional testing (PFT, Level 3), functional performance testing (FPT, level 4), and integrated systems or mission-critical testing (Level 5). PFT ensures individual equipment is correctly installed and operational, while FPT rigorously tests system integration and performance under various conditions, including failure scenarios. These stages confirm the data center is ready for reliable, real-world operations.

In hyperscale data centers, commissioning is conducted at a massive scale and rapid pace. Commissioning a large data center is both expensive and time-sensitive, often accounting for 3–5% of construction costs or $50–$100 per kW. Globally, the data center testing and commissioning market was valued at $9.6 billion in 2023 as reviewed by Data Horizon Research (2025). Despite the high upfront costs, commissioning is seen as a worthwhile investment due to its long-term benefits. Studies show that commissioning typically results in 10–15% energy savings and offers a rapid payback, effectively making it a cost-saving strategy over time (Synergy consulting, 2022).

The commissioning phase in hyperscale data centers is crucial, as delays directly translate to lost revenue and unmet demand. Traditional methods struggle with the complexity of modern systems, which include numerous interconnected subsystems and advanced technologies like AI, IoT, and smart grids. This complexity makes exhaustive testing impractical, increasing the risk of hidden issues. To address these issues and minimize costly delays, the text advocates for automated, data-driven commissioning approaches that can handle the scale and intricacy of today's data centers.

2. Literature review

2.1. Bayesian Optimization in HVAC Commissioning and System Tuning

Machine learning-based HVAC optimization methods like model predictive control (MPC) and reinforcement learning (RL) often demand complex models and extensive data, limiting their scalability. Bayesian Optimization (BO) offers a more cost-effective alternative due to its efficiency with fewer evaluations. Recent research by Fiducioso et al. (2019) has shown BO's effectiveness in HVAC tuning achieved a 32% energy cost reduction by using safe contextual BO for PID controller tuning, while Lu et al. (2020, 2021) demonstrated BO's efficiency in optimizing MPC parameters with reduced computational load. The latter also proposed a reference-model-guided BO to improve convergence by learning system-model residuals.

Bayesian Optimization (BO) has been successfully applied beyond controller auto-tuning to improve design and operational decisions in energy systems integrated by Chakrabarty et al. (2021a). Studies show BO can guide control algorithms, co-optimize system design and control strategies, calibrate HVAC models, and directly control cooling systems Bhattacharya et al. (2021). For example, BO has helped optimize chiller configurations, improved set points in real-time operations, and enhanced energy efficiency by over 10% within weeks of commissioning (Takabatake et al., 2022). While some limitations exist, such as preset adjustment intervals and static safety bounds, overall, BO has proven to be a data-efficient, adaptive tool for optimizing HVAC and cooling systems, making it a strong candidate for AI-driven commissioning (Lin *et al.*, 2023).

Yu *et al.* (2018) formulated a power demand response strategy for distributed data centers using stochastic optimization, with the objective of minimizing total energy cost under uncertain workload and electricity price conditions. This involved evaluating many random demand scenarios to ensure the control decisions (e.g. server dispatch or cooling setpoints) would save cost on average despite the randomness in demand.

Despite such promising results, Bayesian optimization has not yet penetrated mainstream data center commissioning practices. Traditional commissioning largely follows predefined scripts and engineer-driven adjustments, which may not exploit the full potential of data available during the process.

2.2. AI-Driven Commissioning: Current Techniques and Gaps

AI is significantly enhancing data center commissioning through automated fault detection, diagnostics, and predictive analytics. Machine learning analyzes large volumes of sensor data to identify faults and abnormal behavior that might be missed by humans. For example, platforms like Schneider Electric's EcoStruxure have reduced manual checks by 40% during pre-commissioning. Companies like Siemens and Schneider are integrating AI into commissioning to optimize testing and system tuning. In a pilot project in Munich, AI was used to dynamically adjust acceptance criteria based on predicted equipment degradation, potentially cutting maintenance costs by 30%. Overall, AI is transforming commissioning by improving automation, accuracy, and early fault detection. But most AI applications so far have focused on monitoring and analysis (e.g. anomaly detection, simulation) rather than active *decision-making* to optimize the commissioning process itself. In the next section, we outline a stochastic optimization framework that brings Bayesian decision-making into the commissioning process. By doing so, we aim to fill the gap between passive analytics and active optimization, enabling a truly AI-driven commissioning methodology that can handle uncertainty and complex trade-offs automatically.

2.3. Proposed Stochastic Optimization Framework using Bayesian Methods

To modernize data center commissioning, we propose a stochastic optimization framework that orchestrates the commissioning process as an intelligent, closed-loop system. The core of the framework is a Bayesian optimization engine that continuously updates a probabilistic model of the data center's performance as tests are conducted and recommends optimal next actions (tests or adjustments) to achieve commissioning objectives with minimal cost and time. Below, we describe the key elements of the framework, presents a mathematical formulation, and outline the system architecture integrating with on-site commissioning activities.

2.4. Framework Overview and Architecture

Figure 1 illustrates the architecture of the proposed AI-driven commissioning system. The AI-driven commissioning framework for data centers consists of interconnected components designed to optimize testing and system performance through continuous learning and decision-making

2.4.1. Digital Commissioning Model (DCM)

A Bayesian digital twin that starts with a physics-based model and is updated with real-time data. It holds probability distributions of uncertain system parameters and performance metrics, incorporating prior design knowledge and vendor data.

2.4.2. Sensor and Data Integration Layer

Collects and filters high-frequency data from sensors and IoT devices during tests, ensuring clean and reliable inputs for analysis.

2.4.3. Bayesian Optimization Engine

The analytical core that uses surrogate models (e.g., Gaussian Processes) to predict performance, update with new data, and guide the selection of optimal next tests based on an acquisition function.

2.4.4. Commissioning Action Module

Interfaces with human engineers and automation systems to execute and explain AI-suggested test actions, which can be manual or automated.

2.4.5. Bayesian Inference & Learning

Continuously refines the model by updating parameter distributions based on test outcomes, improving predictions and reducing uncertainty iteratively.

2.4.6. Objective & Decision Logic

Frames commissioning as a multi-objective optimization problem, balancing performance compliance, cost/time efficiency, and uncertainty reduction while ensuring safety and regulatory compliance.

This framework enables a data-driven, adaptive, and efficient commissioning process.

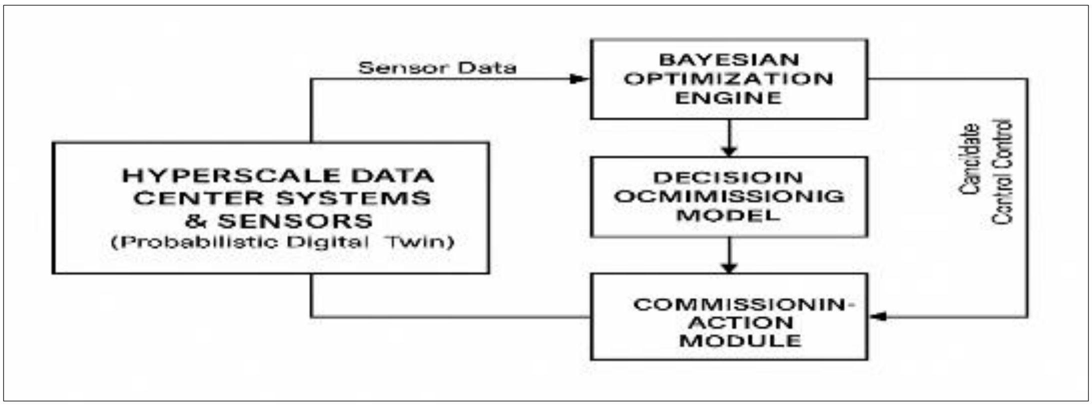

Figure 1 AI-Driven Commissioning Framework

In practice, the commissioning team interacts with this AI-driven system as a decision-support tool. The AI proposes an optimized test plan and sequence, which engineers can review and adjust based on practical considerations. During commissioning execution, the AI monitors results and suggests adaptations – for instance, if a certain test already provides sufficient confidence in a performance area, the AI might suggest skipping a redundant test, saving time. Conversely, if an unexpected outcome occurs, the AI can suggest additional targeted tests to diagnose the issue. The result is a dynamically optimized commissioning process, rather than a static one-size-fits-all checklist.

2.5. Mathematical Formulation

We formalize the commissioning optimization problem as follows. Let **X** represent the space of possible *commissioning actions*. An action $x \in X$ could be a specific test scenario (with certain equipment on/off states, load levels, environmental conditions) or a configuration change (like tuning a parameter). Define **Θ** as a set of uncertain system parameters (e.g. actual UPS efficiency, CRAH unit airflow capacity, sensor biases, etc.) which have a prior distribution *P(Θ)* based on design data and engineering judgment. During commissioning, we cannot directly observe Θ, but we can observe system performance which provides information about Θ.

For each action *x*, define a vector of performance metrics $Y(x) = [y_{1x}, \ y_{2x}), \dots, y_{m(x)}]$, which could include outcomes like energy consumption, thermal stability, redundancy response time, or a binary success/failure for certain criteria. We assume there are desired performance targets or thresholds $T = [T_1, \dots, T_m]$ that must be met (these come from the Owner's Project Requirements and design intent). We introduce an indicator function *I(x)* that equals 1 if action *x* (when applied to the system) results in all performance targets being satisfied (i.e. $y_{i(x)} \geq T_i$) for all required metrics, or within allowable ranges), and 0 if not. In a fully successful commissioning, we seek *I(x) = 1* for the final integrated test of the facility.

We also associate a cost *C(x)* with each action *x*. This cost accounts for time, labor, and any risk or disruption. For example, a full load blackout test might have a high cost (due to complexity and risk), whereas a simple sensor calibration test has a lower cost. Our objective is to find an optimal sequence of actions $\{x_1, \ x_2, \dots, x_N\}$ that minimizes total commissioning cost while achieving $I(x_{final}) = 1$ (all criteria met) with high confidence. We face a dual challenge: we need to both *verify* performance and *optimize* the system (e.g. tune setpoints for efficiency) without excessive testing.

This can be framed as a stochastic optimization or sequential decision problem. One way to cast it is to minimize an expected cost objective:

$$J = \sum_{t=1}^{N} C(x_1) + \wedge . P_r \ \{ \ Performance \ targets \ not \ met \ after \ x_N \},$$

Where the second term is a penalty for failing to meet any performance target by the end of commissioning, with a large weight Λ representing the high cost of an uncommissioned or underperforming facility. The optimization problem is then:

Minimize $E[J]$ over the choice of actions $x_1, \dots, x_N$,

subject to safety constraints for all intermediate steps (each x_t must not violate safe operating limits).

This formulation requires balancing the cost of more testing versus the risk of leaving performance problems unresolved. Because the system's behavior is uncertain (due to unknown Θ), we use a Bayesian approach: we maintain a posterior belief $P_{t(\Theta)}$ after t tests. The probability of meeting all performance targets after a sequence of tests can be evaluated via this belief. We seek to choose each next action $x_{\{t+1\}}$ such that it maximally reduces the expected posterior risk or improves the expected performance.

In a more classic Bayesian optimization framing, one could define a single aggregate objective function $f(x, \theta)$ that measures the "goodness" of the system under configuration or test x and parameters θ. For example, f could be a weighted sum that penalizes unmet criteria heavily and also includes operational efficiency. The true function f is initially unknown (since θ is unknown). We can query f by conducting action f and observing outcomes (with some noise). The goal would be to find the optimal x^* that maximizes f (i.e. yields a fully functional, optimized system) *or* to learn the set of all x that satisfy $I(x) = 1$. We employ Gaussian Process regression to model $f(x)$ as a function with uncertainty. The GP prior captures our initial uncertainty over system responses. Each test provides a data point $(x_t, Y(x_t))$, which updates the GP posterior. We then use an acquisition function $(x; P_t)$ to select the next test $x_{\{t+1\}}$. For commissioning, a suitable acquisition strategy might be a form of Expected Improvement (EI) modified to account for constraints: we seek to maximize the expected improvement in a composite score (e.g. distance to meeting all targets, or expected energy efficiency gain) while ensuring we do not violate known constraints with unsafe tests (Safe BO approaches can ensure we only explore within safe operating bounds (Fiducioso, 2019).

Concretely, at each iteration t we might solve

$$x_{t+1} = \arg\max_{x \epsilon X} \alpha\,(x;\ \ P_t(\Theta))$$

Where $\alpha(x)$ could be, for example, the expected reduction in the probability of any unmet criterion after conducting test x, minus a weighted cost term. This drives the process to focus on tests that either (a) have a high chance to directly achieve commissioning closure (exploitation of a likely good configuration) or (b) greatly inform the uncertainties (exploration where we aren't sure if criteria would pass or fail). Over the course of N selected actions, the process homed in on configurations that meet all requirements and on a calibrated model of the data center's performance.

2.6. Illustrative Use Case in Commissioning Phases

The framework described enhances data center commissioning in two key phases

2.6.1. Pre-Functional Testing (PFT) Optimization

Instead of merely checking if components like CRAH units turn on, the Bayesian Optimization (BO) engine models each unit's performance and recommends specific, informative test points. This approach adapts testing based on results, focusing more on units showing irregularities while minimizing unnecessary tests for well-performing ones. The outcome is a set of calibrated performance models and an initial optimal configuration for energy-efficient operation.

2.6.2. Functional Performance Testing (FPT) Optimization

For system-wide tests (e.g., during simulated failures), the framework intelligently selects the most critical and informative scenarios using BO, rather than relying on limited, manually chosen cases. It identifies potential weak points, suggests tests to validate resilience, and recommends corrective actions if issues are found. This results in more efficient, thorough testing and high statistical confidence that the system will meet performance standards under varied real-world conditions.

2.7. Ensuring Safety and Compliance

A crucial aspect of this Bayesian approach is maintaining safety. Commissioning often involves pushing systems to their limits; doing this in an AI-driven way requires constraints. We integrate "safe optimization" techniques of Fiducioso

(2019), so that the BO engine only suggests tests that are within known safe bounds. For example, the engine would not suggest overloading a generator beyond its rated capacity (which could cause damage), because that constraint is encoded. Instead, if exploring that region is important (e.g. to know the margin), the framework might extrapolate using the probabilistic model or conduct the test in simulation via the digital twin first. Compliance with standards is also monitored – e.g. if a commissioning protocol or code requires a 96-hour continuous run test (Highjoule, 2023), the framework will include that as a mandatory action (but it might still optimize when to schedule it and what data to collect during it). The AI can also handle documentation: each test result is logged and mapped to the required commissioning documentation (generating reports, trend logs, etc. automatically), and even blockchain-like verification for audit trails as suggested in next-gen commissioning proposal can be incorporated.

2.8. Efficiency Improvements in Commissioning Execution

The most direct impact is on the commissioning timeline and labor effort. Traditional commissioning can span many weeks for a large data center, involving iterative testing, manual data review, and sometimes re-testing after fixing issues. Our framework's ability to intelligently select test scenarios can cut down the number of tests required to achieve coverage. For example, instead of blindly performing 10 different load bank tests, the BO engine might determine that 4 well-chosen tests are sufficient to characterize system performance under all critical conditions. This targeted approach can compress the commissioning schedule. Digital twin technology alone has been reported by Polvado (2023) to reduce commissioning time by up to 30% by enabling parallel virtual testing. By adding Bayesian optimization on top, we anticipate even greater reductions, since the AI actively avoids redundant or low-value tests.

Anecdotal evidence from related applications supports this expectation. In the Munich substation case with AI-driven acceptance testing, incorporating predictive algorithms in commissioning reportedly cut the program duration from 14 weeks to just 19 days in a retrofit project. That is nearly 75% reduction in time, attributed to dynamic test criteria and early detection of issues. (Highjoule, 2023). While that example is dramatic and in a specific context, it underlines the potential: by finding problems faster and adapting the plan, AI can eliminate lengthy back-and-forth manual processes. Moreover, automation of data analysis means engineers spend less time analyzing logs and more time on resolving issues, further speeding up the process.

Labor savings translate to cost savings. Fewer test days and more efficient use of expert time directly reduce commissioning contractor fees and internal labor costs. If we conservatively estimate a 20–30% reduction in required commissioning hours (in line with time savings reported with digital tool analyzed by Polvado et al. (2023), that could save hundreds of thousands of dollars on a large project. Additionally, optimization reduces costly repeat tests and delays. In many projects, if a test fails or yields ambiguous results, it must be repeated after adjustments – incurring schedule slips and extra cost. The Bayesian approach, with its focus on information gain, can design tests to maximize clarity, so we get it right the first time. It also quantifies confidence, so we know when we have tested enough. This avoids the tendency to "over-test" out of caution, which sometimes happens in high-stakes facilities. In sum, the commissioning process becomes leaner: doing *just enough* testing to guarantee performance, and no more. Given that commissioning costs for new data centers are on the order of 0.5–1% of capital cost even a one-third reduction represents a notable cost avoidance (for a $200M hyperscale build, 1% would be $2M, so saving 30% of that is ~$600k). This is consistent with experiences from advanced control tuning: Fiducioso et al. (2019) reported a 32% cost reduction in an HVAC control commissioning task using Bayesian optimization**.** We expect similar or greater magnitudes when scaling such approaches to entire data center systems.

2.9. Optimized Performance and Lifecycle Cost Savings

AI-optimized commissioning improves a facility's performance and delivers long-term benefits by automatically fine-tuning system settings during commissioning, rather than relying on default values or manual adjustments later. This data-driven approach enables data centers to start operating near optimal efficiency, reducing energy use and costs over time. For example, Google used AI post-commissioning to cut cooling energy by 40% (siemens 2021), and applying such optimization during commissioning could yield even better initial efficiency. Even small improvements at this stage can lead to substantial savings over a facility's lifetime. (Evans & Gao, 2016).

A key advantage of the framework is enhanced reliability and reduced risk. By using Bayesian-guided testing to explore edge cases, it can uncover hidden defects that traditional methods often miss. For example, the HuiJue Group saw an 82% drop in post-commissioning defect recurrence with AI-informed testing. This leads to fewer operational failures, higher uptime, and avoidance of costly outages, especially critical in data centers, where downtime can cost $10,000 to $30,000 per minute or more. The framework's ability to simulate and test failure scenarios according to Orr et al, (2015) ensures preparedness and supports a proactive, continuous risk management approach throughout the data center's lifecycle

AI-driven commissioning supports long-term benefits by enabling predictive maintenance and continuous commissioning. The Bayesian model created during initial commissioning evolves into a living digital twin, updated with operational data for ongoing optimization. This supports monitoring-based commissioning (MBCx), where the system is continually re-tuned to adapt to changing conditions, preventing performance decline and energy waste. Studies show this approach can deliver an extra 5–10% in energy savings and reduce unplanned downtime. In data centers, integrating AI for both commissioning and operations can cut annual operating expenses by 20% or more by reducing manual intervention and improving efficiency (Compass Datacenters & Schneider Electric, 2025).

Quantitatively, AI-enhanced commissioning significantly reduces risk. For example, lowering the chance of a costly delay from 5% to under 1% potentially saving hundreds of thousands per project. When combined with efficiency gains, it presents a strong business case. Aligning with the industry's shift toward smarter data centers, such as Compass and Schneider's reported 20% OPEX reduction, the framework applies AI early in the lifecycle to enable faster deployment and more efficient, cost-effective operations from the outset.

3. Conclusion

Commissioning is vital for linking a data center’s design with its operational performance. In large-scale (hyperscale) data centers, traditional commissioning methods are no longer sufficient due to their limitations in speed, cost, and reliability. The proposed solution is a Bayesian-based stochastic optimization framework that uses AI to actively guide and improve the commissioning process. This includes real-time data use, and adaptive testing to optimize performance and ensure compliance, all while reducing cost and time.

The framework promises faster deployment, better performance assurance, and long-term operational efficiency, with potential energy savings of 15% and commissioning cost reductions of 30%. It enhances, rather than replaces, human expertise by automating routine tasks and supporting decision-making. The approach aligns with modern industry standards emphasizing data-driven, performance-based practices.

In conclusion, as data centers continue to grow in importance and size (driven by cloud computing and AI workloads), the methods used to commission them must evolve. The stochastic Bayesian optimization framework presented here is a novel contribution toward that evolution. It combines concepts from machine learning, optimization, and building engineering to create a commissioning process that is smarter, faster, and more reliable. Future work will involve prototyping this framework on actual commissioning projects – possibly starting with a specific subsystem (like cooling plants) – to empirically validate the theoretical benefits. We anticipate that as AI-driven commissioning proves its value in reducing lifecycle costs and ensuring quality, it will become a standard practice in hyperscale developments.

Compliance with ethical standards

Disclosure of conflict of interest

No conflict of interest to be disclosed.

References

[1] Bhattacharya, A., Vasisht, S., Adetola, V., Huang, S., Sharma, H., & Vrabie, D. L. (2021). Control co-design of commercial building chiller plant using Bayesian optimization. Energy and Buildings, 246, 111077.

[2] Brown, C., Orr, R., & Rafter, E. (2015). Improve project success through mission critical commissioning. Uptime Institute Journal.

[3] Chakrabarty, A., Danielson, C., Bortoff, S. A., & Laughman, C. R. (2021a). Accelerating self-optimization control of refrigerant cycles with Bayesian optimization and adaptive moment estimation. Applied Thermal Engineering, 197, 117335

[4] Chakrabarty, A., Maddalena, E., Qiao, H., & Laughman, C. (2021b). Scalable Bayesian optimization for model calibration: Case study on coupled building and HVAC dynamics. Energy and Buildings, 253, 111460.

[5] Chen, K.-Y., Li, H., Liu, X., Wang, Z., & Song, Y. (2021). Automated Commissioning of Data Centers with an Intelligent Bayesian Optimization Engine. 2021 IEEE International Conference on Big Data

[6] Compass Datacenters & Schneider Electric. (2025). Compass and Schneider Electric utilize AI to transform data center maintenance (Press Release).

[7] Data Horizon Research. (2025). Data Center Testing and Commissioning Service Market Size, Growth, Share & Analysis Report – 2033. (Market research report).

[8] Dataknox (Lamei, S.). (2025). 5 stages of data center commissioning [Blog post].

[9] Evans, R., & Gao, J. (2016). DeepMind AI reduces Google data centre cooling bill by 40%. DeepMind Blog. Retrieved from deepmind.google blog.

[10] Fiducioso, M., Curi, S., Schumacher, B., Gwerder, M., & Krause, A. (2019). Safe contextual Bayesian optimization for sustainable room temperature PID control tuning. In Proceedings of IJCAI 2019. (arXiv:1906.12086).

[11] Highjoule (2023) Innovative energy storage solutions for home, industrial & commercial use - HighJoule. https://www.highjoule.com/products/ (Accessed: June 7, 2025).

[12] Lin, X., Guo, Q., Yuan, D., & Gao, M. (2023). Bayesian optimization framework for HVAC system control. Buildings, 13(2), 314.

[13] Lu, Q., González, L. D., Kumar, R., & Zavala, V. M. (2021). Bayesian optimization with reference models: A case study in MPC for HVAC central plants. Computers & Chemical Engineering, 154, 107491.

[14] Lu, Q., Kumar, R., & Zavala, V. M. (2020). MPC controller tuning using Bayesian optimization techniques. arXiv preprint, arXiv:2009.14175.

[15] Mark H., Seidi R, Shalley C (2005), Data Centers: Staying on-line, data center commissioning. Ashare journal. 18(2), 25–39

[16] Polvado, A. (2023). Efficient data center commissioning: 6 ways digital twins are game changers. LinkedIn Pulse. Retrieved from linkedin.com

[17] Siemens. (2021). How artificial intelligence is cooling data center operations (White paper). Siemens AG, USA.

[18] Synergy Consulting Engineers (2022) What is the ROI of Commissioning? https://synergy-engineers.com/what-is-the-roi-of-commissioning (Accessed: June 7, 2025).

[19] Takabatake, T., Yamamoto, M., & Hino, H. (2022). Algorithm for searching optimal set values of an absorption chiller system using Bayesian optimization. Science and Technology for the Built Environment, 28(2), 188–199.

[20] Uptime Institute. (2024a). 2024 global data center survey – Executive summary. Uptime Institute Intelligence Report.

[21] Uptime Institute. (2024b). 2024 AI and software in data centers survey – Executive summary. Uptime Institute Intelligence Report.

[22] Wang, Z., Lin, X., Zhang, H., et al. (2024). Operation optimization considering multiple uncertainties for the multi-energy system of data center parks based on information gap decision theory. Frontiers in Energy Research, 12, 1423126.

[23] Yu, X., Xue, Y., & Li, L. (2018). Real-time energy management of large-scale data centers: A stochastic approach. IEEE Transactions on Sustainable Computing, 3(3), 226–239.