








Original research article

Smart HVAC Heat Exchanger Network Optimization Through Collaborative IoT-Enabled Predictive Analytics for Manufacturing Facilities

F. Bakhritdinov^a  0009-0001-3684-8754, Z. Atamuratova^{b,c}  0009-0006-2774-2612,
S. Sabirov^d  0009-0008-0504-7568, A. Umarov^e  0000-0003-2408-3624,
A. M. Alsayah^{f,*}  0009-0005-8122-8182

^a *Kimyo International University in Tashkent, Shota Rustaveli str. 156, Tashkent 100121, Uzbekistan;*

^b *National Research University TIIAME, Kori Niyoziy 39, Tashkent 100000, Uzbekistan;*

^c *Urgench State University, Kh. Alimdjani str. 14, Urgench 220100, Uzbekistan;*

^d *Mamun University, Bolkhovuz Street 2, Khiva 220900, Uzbekistan;*

^e *University of Tashkent for Applied Sciences, Str. Gavhar 1, Tashkent 100149, Uzbekistan;*

^f *Refrigeration & Air-condition Department, Technical Engineering College, The Islamic University, Najaf, Iraq*

ABSTRACT

Manufacturing facilities rely on sophisticated Heating, Ventilation, and Air Conditioning (HVAC) systems to ensure precise environmental conditions; however, operating these systems in isolated silos often results in substantial energy inefficiencies. This study addresses this challenge by developing and validating a collaborative Internet of Things enabled framework that optimizes heat exchanger networks using privacy-preserving predictive analytics. A distributed IoT architecture comprising 1,234 sensors was deployed across eight diverse manufacturing facilities (chemical, electronics, and automotive) in Saudi Arabia. The framework utilized federated Long Short-Term Memory neural networks. Using the Federated Averaging algorithm, these networks collaboratively trained a global optimization model without sharing proprietary local data. Over a 12-month operational period compared against a three-month baseline, the framework achieved a 29.1% average reduction in HVAC energy consumption ($p < 0.001$) and improved temperature control precision by 37%. Furthermore, the federated learning model significantly outperformed isolated control strategies, reducing prediction error by 61.8% and preventing 94% of inter-zonal operational conflicts. These results demonstrate that collaborative, privacy-preserving intelligence offers a scalable, robust solution for industrial energy management, effectively bridging the gap between localized control and system-wide optimization in support of Industry 5.0 sustainability goals.

ARTICLE INFO

Article history:

Received August 19, 2025

Revised December 5, 2025

Accepted January 23, 2026

Published online April 23, 2026

Keywords:

Collaborative optimization;

Federated learning;

HVAC systems;

IoT integration;

Predictive analytics

*Corresponding author:

Ahmed Mohsin Alsayah

ahmadalsayah@iumajaf.edu.iq

1. Introduction

In the Industry 4.0 era, the imperative for sustainable and energy-efficient operations within manufac-

turing facilities has become a critical focus of industrial and environmental policy [1], [2]. Manufacturing plants, particularly high-technology fabrication facilities, are characterized by substantial energy consumption, with Heating, Ventilation, and Air Conditioning

(HVAC) systems often representing 30-65% of the total energy budget [3]-[5]. These systems are not merely for comfort but are integral to maintaining the stringent environmental conditions—such as precise temperature, humidity, and air quality—necessary for high-yield production and product integrity [6]. In regions with demanding climates, such as the Eastern Province of Saudi Arabia, the energy load attributed to HVAC systems is further exacerbated, making optimization a strategic priority for both economic competitiveness and environmental stewardship [7], [8]. Inefficient operation of these systems leads to significant energy waste, increased operational costs, and a substantial carbon footprint, directly impacting profitability and regulatory compliance [9], [10].

The evolution of HVAC control has progressed from rudimentary localized thermostats to sophisticated, data-centric strategies. Initial advancements focused on improving the efficiency of individual components and implementing rule-based controls [11]. More recently, the integration of the Internet of Things (IoT) has enabled the collection of vast amounts of real-time operational data, paving the way for data-driven optimization [12], [13]. Current state-of-the-art approaches predominantly leverage Machine Learning (ML) and Artificial Intelligence for single-system optimization [14]-[16]. Methodologies such as Model Predictive Control (MPC) have been implemented to optimize energy consumption and thermal comfort by forecasting future needs based on historical data and weather predictions [17]. Studies have demonstrated that occupancy-based controls and smart energy management systems can yield energy savings of approximately 17.8% and 37%, respectively, within a single facility [18]. Furthermore, research into the optimization of Heat Exchanger Networks (HENs), a critical component of industrial HVAC systems, has traditionally relied on pinch analysis and mathematical programming [19], [20]. While effective, these methods often treat the system statically and do not dynamically adapt to changing operational conditions or account for the time-dependent nature of processes like heat exchanger fouling [21]. Modern approaches have begun to apply ML to predict HEN performance and identify optimal configurations, but typically within the confines of a single, isolated process network [22]. Despite these advancements, a significant limitation persists: HVAC systems and their associated HENs typically operate in isolation, lacking the capability to communicate or coordinate with systems in other zones or facilities [22], [23]. This operational siloing prevents system-wide optimization and the sharing of learned

insights, representing a substantial untapped potential for energy savings and performance improvements.

The primary aim of this study is to design, implement, and validate a collaborative IoT-enabled framework that leverages predictive analytics and Federated Learning (FL) to optimize HVAC heat exchanger networks across multiple manufacturing facilities. The specific objectives are:

1. To deploy a distributed IoT sensor network across eight manufacturing facilities in Saudi Arabia to capture real-time data on HVAC performance and environmental parameters.
2. To develop and integrate LSTM-based predictive models to forecast energy demand, operational loads, and potential system conflicts.
3. To implement a FL architecture to enable collaborative training of a global optimization model while ensuring the data privacy of each participating facility.
4. To quantify the framework's effectiveness in reducing energy consumption, enhancing temperature control precision, and lowering operational costs.

The principal contribution of this work is the development and empirical validation of a collaborative FL framework specifically optimized for industrial HENs. Distinct from prior research that relies on single-facility MPC or simulated multi-agent systems, this study implements a physical, privacy-preserving architecture across heterogeneous manufacturing sites. This approach introduces a novel "collaborative uplift" mechanism, demonstrating how stable operational environments can enhance predictive performance for facilities with volatile thermal loads without data exchange. Consequently, this work establishes a scalable, validated pathway for Industry 5.0 energy management that reconciles system-wide optimization with strict data sovereignty requirements.

2. Literature Review

Recent literature from 2024 and 2025 has begun to acknowledge this limitation, shifting focus toward decentralized and collaborative intelligence. Centralized models often require the external aggregation of sensitive operational data. Conversely, isolated systems cannot benefit from cross-facility patterns. FL addresses these limitations by enabling collective optimization while preserving local data governance. For instance, research demonstrated that FL can significantly improve the generalization of autonomous

HVAC control policies across diverse climate zones, outperforming isolated agents [24], [25]. Similarly, Feng [26] highlighted the growing role of ML in optimizing building energy management through parameter tuning and operational variable control. However, these emerging studies primarily rely on simulation environments or single-building scenarios, leaving a gap for empirical validation in large-scale, heterogeneous manufacturing networks where operational constraints are more severe.

A brief overview of seminal and recent studies is presented in Table 1, which compares the methodologies, focus areas, and limitations of existing research in HVAC and HEN optimization. This comparison highlights the prevailing focus on single-instance optimization and underscores the gap this study aims to address.

Synthesizing the methodologies detailed in Table 1, current research can be categorized into localized physical controls, which ensure stability but lack adaptability, and centralized data-driven models, which optimize performance but compromise data privacy. Although distributed learning offers a theoretical middle ground, the literature review reveals a clear and critical gap: the absence of an empirically validated framework that enables collaborative, predictive optimization across multiple, independent

HVAC systems and their heat exchanger networks. While predictive analytics for single-building optimization is an active area of research [34], these systems operate as isolated "islands of intelligence." They do not share operational insights, learned efficiencies, or predictive models with other systems, thereby missing opportunities for global optimization [35], intelligent load balancing, and collective problem-solving. Consequently, each facility is forced to learn independently, a process that is inefficient, redundant, and fails to leverage the collective data and experience of the wider industrial ecosystem. To date, no studies have demonstrated a practical implementation of a decentralized, privacy-preserving learning architecture for the express purpose of coordinating HEN operations across disparate manufacturing facilities.

This study addresses the identified gap by proposing a novel collaborative framework. The choice of a distributed IoT architecture is justified by the need for granular, real-time data acquisition from a wide array of sensors across geographically separate facilities [36], [37]. To overcome the challenge of data siloing and the inherent privacy concerns associated with sharing proprietary operational data, this research employs FL. FL is a decentralized ML approach that enables multiple entities to collaboratively train a shared prediction model without ex-

Table 1. Comparative analysis of recent literature in HVAC and hen optimization

Reference	Methodology	Focus Area	Key Findings	Identified Limitations
[27]	MPC, IoT	Single-zone HVAC energy and comfort optimization	MPC can effectively balance energy use and thermal comfort in a closed-loop system.	Limited to a single zone; does not address inter-system coordination or network effects.
[28], [29]	Mathematical Programming, ML	Heat Exchanger Network (HEN) design and operational optimization	ML and advanced algorithms can identify optimal HEN configurations to reduce energy consumption.	Primarily focused on design-stage or single-process optimization; lacks real-time, dynamic adaptation across multiple systems.
[30]	IoT, Data Analytics, Predictive Maintenance	Single-facility energy efficiency and equipment health	IoT-enabled predictive maintenance reduces downtime and optimizes resource allocation within one facility.	Operates in an information silo; no mechanism for sharing learned failure patterns or optimization strategies with other facilities.
[31], [32]	Centralized AI, Explainable AI (XAI)	Smart building energy management	Centralized data collection and AI analysis can optimize building performance but raises significant privacy and security concerns.	Centralization creates data bottlenecks and privacy risks; not scalable for collaborative optimization among different entities.
[33]	FL Review	Distributed ML for energy systems	FL is a promising paradigm for privacy-preserving collaborative model training in the energy sector.	Primarily a review of potential; lacks application-specific implementation for collaborative HVAC/HEN control in manufacturing.

changing their local data, thus preserving data privacy and security [38]-[41]. This methodology is superior to centralized approaches that require sensitive data to be aggregated in a single location, which introduces significant security risks and is often untenable for competing commercial entities. For predictive analytics, Long Short-Term Memory (LSTM) neural networks are utilized. This choice is justified by the proven efficacy of LSTMs in modeling complex, time-series data, which is characteristic of HVAC energy consumption, weather patterns, and manufacturing schedules [42], [43].

3. Methodology

3.1 Study Design and Experimental Setting

This study was conducted over a 12-month period, from August 2023 to July 2024, to capture a full cycle of seasonal variations. The research employed a multi-site, longitudinal design involving eight distinct manufacturing facilities located within the industrial hubs of Dammam and Jubail in the Eastern Province of Saudi Arabia. The participating facilities were selected to represent a diverse range of manufacturing typologies, including two electronics assembly plants, three chemical processing facilities, two automotive parts manufacturers, and one pharmaceutical production plant. This diversity ensured the robustness and generalizability of the developed framework across different operational demands and HVAC load profiles. Prior to the implementation of the collaborative framework, a three-month baseline period was established to collect operational data under the existing, non-collaborative control systems. This baseline data served as a benchmark for evaluating the performance improvements achieved by the new system. All facilities operated on a 24/7 basis, with variable production schedules that were incorporated as inputs into the predictive models.

3.2 Distributed IoT Architecture and Data Acquisition

A comprehensive IoT architecture was designed and deployed to enable granular, real-time data collection from the HVAC systems and associated HENs at each of the eight facilities. A total of 1,234 industrial-grade sensors were installed. The sensor suite at each facility was standardized and included: high-precision temperature sensors ($\pm 0.1^\circ\text{C}$ accuracy) monitoring shell-and-tube inlet and outlet tempera-

tures; ultrasonic flow meters quantifying fluid mass flow rates; and feedback actuators recording control valve positions (0–100%). Environmental monitoring comprised relative humidity sensors ($\pm 2\%$ RH accuracy), differential pressure sensors for monitoring airflow across filters and coils, and air quality sensors measuring Carbon Dioxide (CO_2) and *Total volatile organic compounds*. Additionally, non-invasive smart energy meters monitored the power consumption of chillers, pumps, and air handling units, while occupancy was tracked using a combination of *passive infrared* sensors and network-based device counting to provide an accurate measure of internal heat loads from personnel.

Data from these sensors were collected at a five-minute sampling interval, generating a high-resolution dataset exceeding 129 million unique data points over the 12-month operational period. This high-frequency acquisition was managed by local IoT gateways (Raspberry Pi 4 Model B units running a custom Linux build). These gateways were responsible for data aggregation, temporary local storage, and secure transmission. For communication, the *Message queuing telemetry transport* protocol was utilized due to its lightweight nature and efficiency in constrained network environments. All data transmission between the sensors, gateways, and the local facility servers was secured using *Transport layer security* v1.3 encryption, with X.509 certificate-based authentication for all devices to ensure data integrity and prevent unauthorized access. Each facility maintained a local time-series database (InfluxDB) to store its own operational data, ensuring data ownership and physical separation.

3.3 Collaborative Optimization Framework

The core of this research is the development of a collaborative optimization framework that integrates predictive modeling, federated learning, and real-time control. This framework was designed to operate in a decentralized manner, allowing for shared intelligence without the centralized collection of sensitive operational data.

3.4 Predictive Modeling with Long Short-Term Memory Networks

To forecast future HVAC energy demands and internal thermal loads, an LSTM neural network was employed. LSTMs are a type of *recurrent neural network* particularly well-suited for time-series forecasting due to their ability to learn long-term depen-

dencies in data. The LSTM model at each facility was trained to predict the next 24 hours of energy consumption and zone temperatures based on a 168-hour (7-day) look-back period of historical data.

The input vector x_t for the LSTM at time t comprised the following features: historical sensor readings (temperature, humidity, energy consumption), local weather forecast data (ambient temperature, solar irradiance) obtained from a public API, and the facility's production schedule (operational hours, machine usage). The core of the LSTM cell is governed by a series of gates that control information flow. The forget gate, f_t , determines which information from the previous cell state, C_{t-1} , should be discarded. Its operation is described by Equation (1), where σ is the sigmoid function, W_f and U_f are weight matrices, h_{t-1} is the previous hidden state, and b_f is the bias vector [9], [44].

$$f_t = \sigma(W_f x_t + U_f h_{t-1} + b_f) \quad (1)$$

The input gate, i_t , shown in Equation (2), decides which new information will be stored in the cell state. A candidate cell state, \tilde{C}_t , is created using a \tanh function as shown in Equation (3) [16].

$$i_t = \sigma(W_i x_t + U_i h_{t-1} + b_i) \quad (2)$$

$$\tilde{C}_t = \tanh(W_c x_t + U_c h_{t-1} + b_c) \quad (3)$$

The new cell state, C_t , is then updated by combining the forget gate's output with the old state and the input gate's output with the candidate state, as formalized in Equation (4). This mechanism allows the model to selectively remember or forget patterns over long durations, which is crucial for predicting HVAC loads influenced by weekly or seasonal cycles [42].

$$C_t = f_t \otimes C_{t-1} + i_t \otimes \tilde{C}_t \quad (4)$$

Finally, the output gate, o_t , and the new hidden state, h_t , are calculated according to Equations (5) and (6), producing the prediction for the next time step [29], [45].

$$o_t = \sigma(W_o x_t + U_o h_{t-1} + b_o) \quad (5)$$

$$h_t = o_t \otimes \tanh(C_t) \quad (6)$$

The model was implemented in Python using the TensorFlow and Keras libraries. The network architecture consisted of two stacked LSTM layers with 128 units each, followed by a dense output layer.

3.5 Privacy-Preserving Federated Learning Architecture

To enable collaborative learning without sharing raw data, an FL architecture was implemented using the FedAvg algorithm. This approach allowed the individual LSTM models at each of the K facilities to benefit from the experiences of all other facilities, leading to a more robust and accurate global model. The process was coordinated by a central server that did not have access to any facility's data.

The iterative process, schematically illustrated in Figure 1, proceeded in communication rounds. At the beginning of a round t , the central server distributes the current global model weights, w_t , to a subset of k participating facilities. Each selected facility k then trains the model on its local data for a number of epochs, computing an updated set of local weights, w_{t+1}^k , by minimizing its local loss function. After local training, each facility sends only its updated model weights back to the server. The central server then aggregates these weights to produce a new global model for the next round, w_{t+1} . The aggregation is performed using a weighted average, as shown in Equation (7), where n_k is the number of data samples at facility k , and N is the total number of data samples across all participating facilities [4].

$$w_{t+1} = \sum_{k=1}^K \frac{n_k}{N} w_{t+1}^k \quad (7)$$

This process was repeated for 200 communication rounds, with 5 local training epochs per round at each facility. This federated approach ensures that proprietary operational data never leaves the facility's local network, thus preserving data privacy and security while still enabling the creation of a collectively intelligent predictive model.

3.6 Real-Time Distributed Control Algorithm

The predictions generated by the federated LSTM model were fed into a real-time, multi-objective optimization algorithm deployed on the local controllers at each facility. The objective of this algorithm was to determine the optimal operational setpoints for the HVAC and HEN components (e.g., chiller output, air handling unit fan speeds, valve positions) to minimize a composite cost function. The cost function, J , as defined in Equation (8), is a weighted sum of three competing objectives: energy consumption cost (J_{energy}), thermal comfort deviation penalty (J_{comfort}), and peak demand cost (J_{peak}) [28].

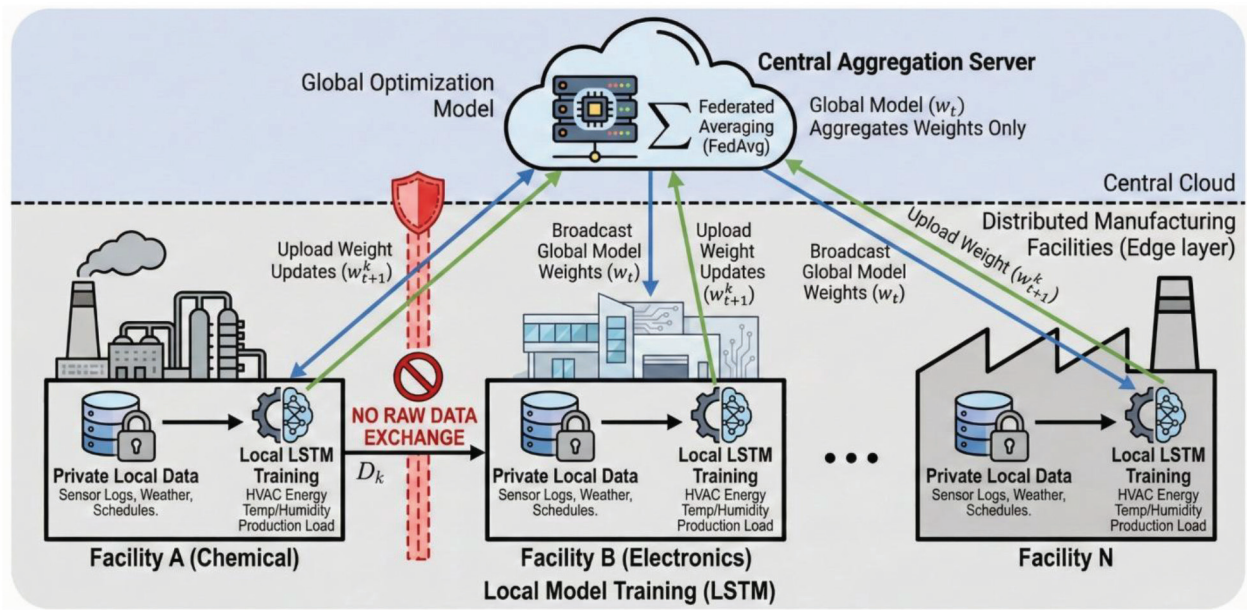


Figure 1. Schematic of the collaborative FL architecture. The central server initializes and distributes the global model to local facilities. Each facility trains the model on private, on-premise data and returns only the weight updates. The server aggregates these updates using the FedAvg algorithm to refine the global model, which is then redistributed, ensuring data privacy is maintained throughout the cycle.

$$\min J = \alpha_1 J_{\text{energy}} + \alpha_2 J_{\text{comfort}} + \alpha_3 J_{\text{peak}} \quad (8)$$

Here, α_1 , α_2 , α_3 are weighting coefficients that were tuned to balance the trade-offs between energy savings, occupant comfort, and grid stability. J_{energy} was calculated based on the predicted power consumption and real-time electricity tariffs. The thermal comfort penalty, J_{comfort} , was formulated as a quadratic function that penalizes deviations from the desired temperature setpoint range ($T_{\text{setpoint}} \pm 1^\circ\text{C}$), ensuring precise environmental control. J_{peak} was designed to penalize high power consumption during peak tariff periods, encouraging load shifting. The optimization problem was solved at each 15-minute control interval using a *particle swarm optimization* algorithm, which is well-suited for complex, non-linear problems.

3.7 Performance Metrics and Evaluation

To rigorously evaluate the effectiveness of the implemented framework, a set of Key Performance Indicators (KPIs) was established. The performance was compared against the data collected during the three-month baseline period.

The primary metric was the *Energy Consumption Reduction*, calculated as the percentage decrease in total kilowatt-hours (kWh) used by the HVAC systems. This was determined using Equation (9) [42], [46]:

$$\text{EnergySavings}(\%) = \left(1 - \frac{\text{Energy}_{\text{post}}}{\text{Energy}_{\text{baseline}}} \right) \times 100 \quad (9)$$

Temperature Control Precision was quantified by the reduction in the standard deviation of the measured indoor air temperature from the defined setpoint. A lower standard deviation indicates more stable and precise control. The improvement was calculated as shown in Equation (10) [27]:

$$\text{PrecisionImprovement}(\%) = \left(1 - \frac{\sigma_{\text{post}}}{\sigma_{\text{baseline}}} \right) \times 100 \quad (10)$$

where σ_{post} and σ_{baseline} are the standard deviations of the temperature during the post-implementation and baseline periods, respectively.

Other metrics included *Peak Demand Reduction*, measured in kilowatts (kW); the number of *Prevented System Conflicts*, defined as instances of simultaneous heating and cooling in adjacent, interlinked zones; and *Aggregate Energy Cost Savings*, calculated in U.S. Dollars based on the prevailing industrial electricity tariffs in Saudi Arabia.

3.8 Statistical Analysis

To validate the statistical significance of the observed performance improvements, this study employed a rigorous hypothesis testing framework using

Python's *statsmodels* and *scipy.stats* libraries. Given the longitudinal nature of the data, paired t-tests were selected to compare daily energy consumption and temperature deviations between the baseline and post-implementation phases for each facility, effectively controlling for site-specific fixed effects. Prior to testing, the normality of the paired difference distributions was verified using the Shapiro-Wilk test to ensure the validity of the parametric approach. A significance threshold (α) of 0.05 was established, and p-values were calculated to quantify the probability that observed reductions in energy use and control error were attributable to the framework's intervention rather than random variation. Finally, correlation analyses were conducted to examine the relationship between the model's predictive accuracy and the realized magnitude of energy savings.

3.9 Model Validation Strategy

Beyond these statistical tests, this study implemented a specific validation protocol for the federated LSTM model prior to real-time deployment. To account for the temporal dependencies inherent in HVAC data, a time-series cross-validation (sliding window) technique was employed instead of random k-fold validation. The historical dataset for each facility was partitioned chronologically, designating the first 80% for training and the subsequent 20% for validation, ensuring that the model was tested strictly on unseen future data to prevent look-ahead bias. The model hyperparameters were tuned to minimize the validation loss function. This rigorous offline validation phase confirmed the model's stability and generalization capability before it was permitted to generate control signals in the live manufacturing environment.

4. Results and Discussions

The implementation of the collaborative IoT-enabled framework yielded substantial and statistically significant improvements across all monitored performance domains. The results are presented in a sequence that first establishes the predictive accuracy of the core federated model, then quantifies the resultant efficiencies in energy consumption and operational control and finally explores the dynamics of the collaborative learning process itself. All findings reported herein are derived from the 12-month operational period and are compared against the initial three-month baseline period.

4.1 Performance and Accuracy of the Federated Predictive Model

The efficacy of the entire optimization framework is predicated on the accuracy of its underlying predictive model. Therefore, a comprehensive analysis was conducted to evaluate the performance of the federated LSTM network. The analysis focused on the model's ability to accurately forecast key operational variables, the tangible benefits derived from the FL approach compared to isolated models, and the convergence characteristics of the collaborative training process. The results of this multi-faceted evaluation are presented in Figure 2, which is structured across four analytical panels to provide a complete picture of model performance. Figure 2a illustrates the model's dynamic tracking capability by plotting the predicted versus actual total HVAC energy consumption over a representative one-week period for a single facility. Figure 2b provides a quantitative assessment of the model's accuracy in predicting zone temperatures through a correlation plot. To substantiate the value of the collaborative methodology, Figure 2c presents a comparative analysis of prediction error between the federated model and standalone, locally trained models for all eight facilities. Finally, Figure 2d visualizes the learning process by charting the reduction in the global model's validation loss across the 200 communication rounds of federated training.

The federated LSTM model demonstrated a high degree of accuracy in forecasting operational parameters. As shown in Figure 2a, the model's 24-hour-ahead prediction of total HVAC energy consumption closely tracked the actual measured values. The model effectively captured complex daily fluctuations driven by production schedules, ambient weather conditions, and internal thermal loads. The model adeptly anticipated both the sharp ramps in energy use corresponding to the start of production shifts and the subsequent reductions during lower-activity periods. Quantitative analysis of temperature prediction, presented in Figure 2b, further confirms the model's precision. The scatter plot visually demonstrates a tight clustering of data points along the 45-degree identity line, illustrating the strong agreement between the forecasted and measured values. This linear alignment corresponds to a coefficient of determination (R^2) of 0.98, indicating that the model captured 98% of the variance in thermal dynamics. As shown by the bounds in the plot, the vast majority of predictions fell within the narrow $\pm 1.0^\circ\text{C}$ error margin, a critical requirement for maintaining the stringent environmental conditions necessary in the participating manufacturing facilities [13].

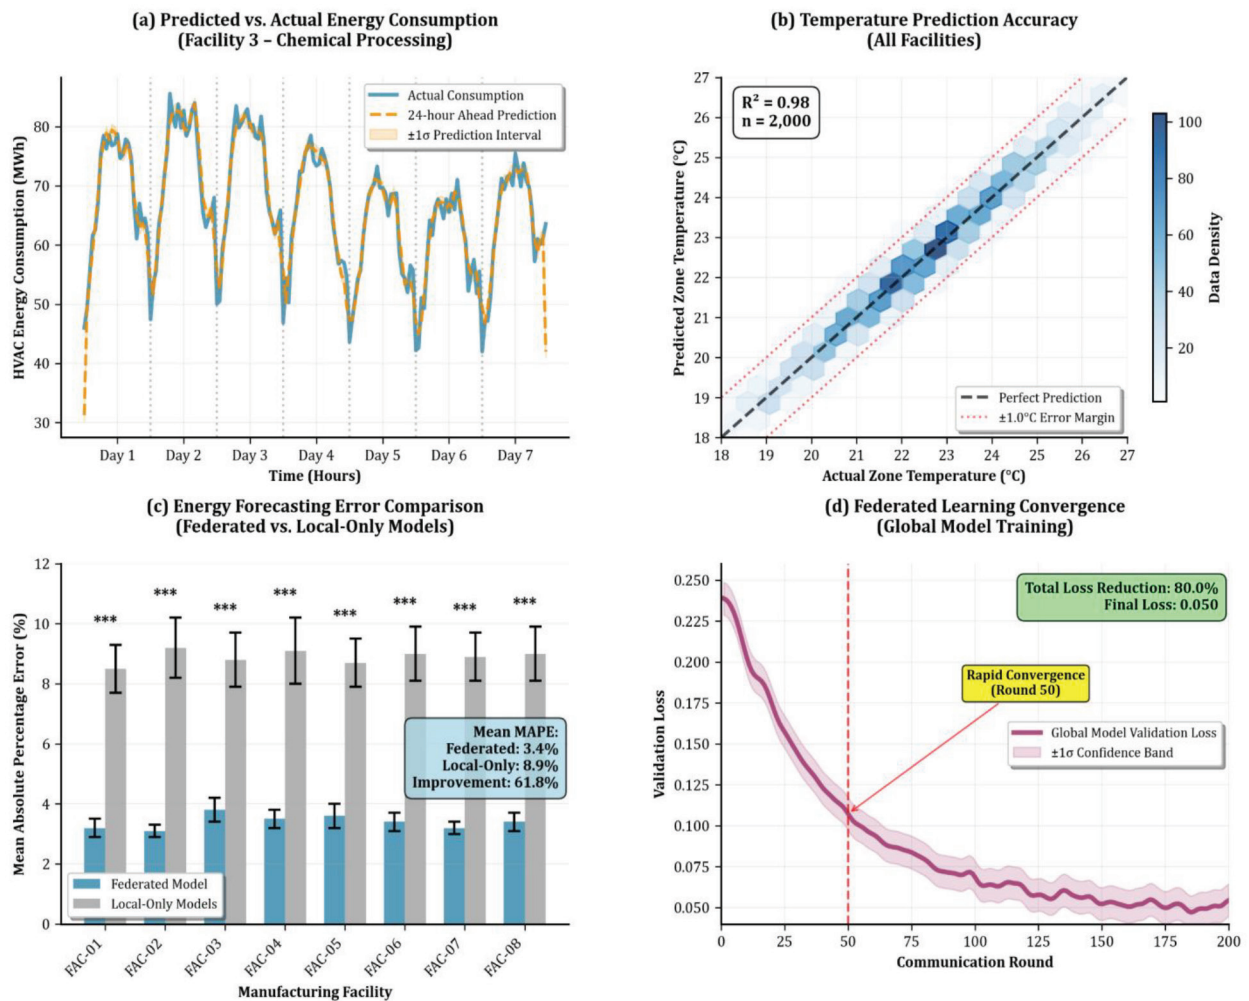


Figure 2. Federated LSTM Model Performance and Convergence. (a) Predicted versus actual HVAC energy consumption over one week. (b) Correlation of predicted versus actual zone temperatures, with an R^2 of 0.98. (c) MAPE in energy forecasting for the federated model compared to local-only models. (d) FL convergence curve showing validation loss reduction over 200 communication rounds.

A central hypothesis of this study was that a collaborative, FL approach would outperform isolated, single-facility models. The comparative bar chart in Figure 2c provides clear evidence supporting this hypothesis, displaying a consistent performance gap between the two approaches. The federated model achieved a significantly lower Mean Absolute Percentage Error (MAPE) in energy forecasting across all eight facilities when compared to the taller bars representing the local-only models. Quantitatively, the average MAPE for the federated model was 3.4%, whereas the average for the local-only models was 8.9%. This difference represents a 61.8% reduction in prediction error, a direct result of the model's ability to learn from the more diverse dataset provided by the entire network. The learning dynamics are further detailed in Figure 2d, where the validation loss curve exhibits a sharp, monotonic descent within the first 50 communication rounds before leveling

off. This trajectory indicates that the global model converged rapidly to a stable minimum, confirming that the FedAvg algorithm effectively aggregated local insights without signs of instability or divergence [15].

4.2 Framework Efficacy in Energy Consumption and Cost Reduction

The primary objective of the collaborative framework was to translate predictive accuracy into tangible reductions in energy consumption, peak demand, and operational expenditure. The implementation of the real-time distributed control algorithm, informed by the federated LSTM's predictions, led to substantial improvements in these KPIs. A detailed summary of the performance outcomes for each of the eight participating facilities is provided in Table 2. The table presents a comparative analysis of key metrics from the three-month baseline period versus the 12-month

post-implementation period, including mean daily energy consumption, percentage of energy savings, reduction in peak grid demand, and the resulting annualized cost savings. All observed reductions in energy consumption were found to be statistically significant.

As detailed in Table 2, the collaborative framework achieved an average energy consumption reduction of 29.1% across all participating facilities, lowering the total daily consumption from 490,850 kWh to 344,258 kWh. The data reveals distinct performance tiers based on facility type; chemical processing plants (FAC-03, FAC-04, FAC-05) consistently achieved the highest savings (mean 32.3%) and financial returns, driven by their continuous, high-inertia thermal loads which maximize the utility of predictive pre-cooling. In contrast, discrete manufacturing sites like electronics assembly (FAC-01, FAC-02) and pharmaceutical production (FAC-08) realized slightly lower yet significant savings (24.0% to 27.0%), reflecting the constraints imposed by variable shift patterns and stricter instantaneous environmental tolerances. Statistical validation via paired t-tests confirms that these reductions are significant ($p < 0.001$) for every individual facility, validating the framework's efficacy across heterogeneous operational profiles.

In addition to overall energy reduction, the framework was highly effective at managing peak electricity demand. By predictively pre-cooling zones and intelligently coordinating the operation of high-load equipment across the network, the system achieved an average peak demand reduction of 283.8 kW per facility. This intelligent load-shifting capability is critical for reducing demand charges, which constitute a significant portion of industrial electricity bills [23]. The combined effect of reduced energy consump-

tion and lower peak demand resulted in substantial financial savings, totaling an aggregate of \$2,120,980 annually. As detailed in Table 2, the chemical processing facilities contributed the largest share of these savings (approximately 50%, averaging \$353,983 per facility), driven by their high energy intensity and flexibility for load shifting. Automotive parts manufacturers followed with an average saving of \$267,350 per facility, while the electronics and pharmaceutical plants, constrained by tighter environmental tolerances, achieved lower but significant average savings of \$174,777 per facility.

4.3 Enhancement of Environmental Control and Operational Stability

Beyond energy efficiency, the framework was designed to improve the precision of environmental control and the stability of HVAC operations. The system's performance in these areas was evaluated by analyzing the deviation of zone temperatures from their required setpoints and by quantifying the reduction in system conflict events, such as simultaneous heating and cooling of adjacent zones. The results of this analysis are presented in Figure 3. The figure uses a two-panel structure to address these two distinct aspects of operational performance. Figure 3a uses violin plots to compare the distribution of temperature deviations from the setpoint during the baseline and post-implementation periods, providing a visual representation of control precision. Figure 3b employs a bar chart to show the dramatic reduction in the monthly count of system conflict events, illustrating the improvement in operational logic and coordination.

Table 2. Summary of energy consumption, peak demand, and cost savings by facility

Facility ID	Facility Type	Mean Daily Baseline Consumption (kWh)	Mean Daily Post-Implementation Consumption (kWh)	Energy Savings (%)	Peak Demand Reduction (kW)	Annualized Cost Savings (USD)	p-value (Paired t-test)
FAC-01	Electronics Assembly	48,550	35,442	27.0%	215	\$188,750	<0.001
FAC-02	Electronics Assembly	51,200	37,376	27.0%	230	\$201,360	<0.001
FAC-03	Chemical Processing	75,300	51,204	32.0%	350	\$348,520	<0.001
FAC-04	Chemical Processing	79,800	54,264	32.0%	385	\$369,150	<0.001
FAC-05	Chemical Processing	72,100	48,307	33.0%	340	\$344,280	<0.001
FAC-06	Automotive Parts	63,400	44,380	30.0%	290	\$275,090	<0.001
FAC-07	Automotive Parts	61,900	43,949	29.0%	280	\$259,610	<0.001
FAC-08	Pharmaceutical	38,600	29,336	24.0%	180	\$134,220	<0.001
Mean	–	61,356	43,032	29.1%	283.8	\$265,122	–
Total	–	490,850	344,258	–	2,270	\$2,120,980	–

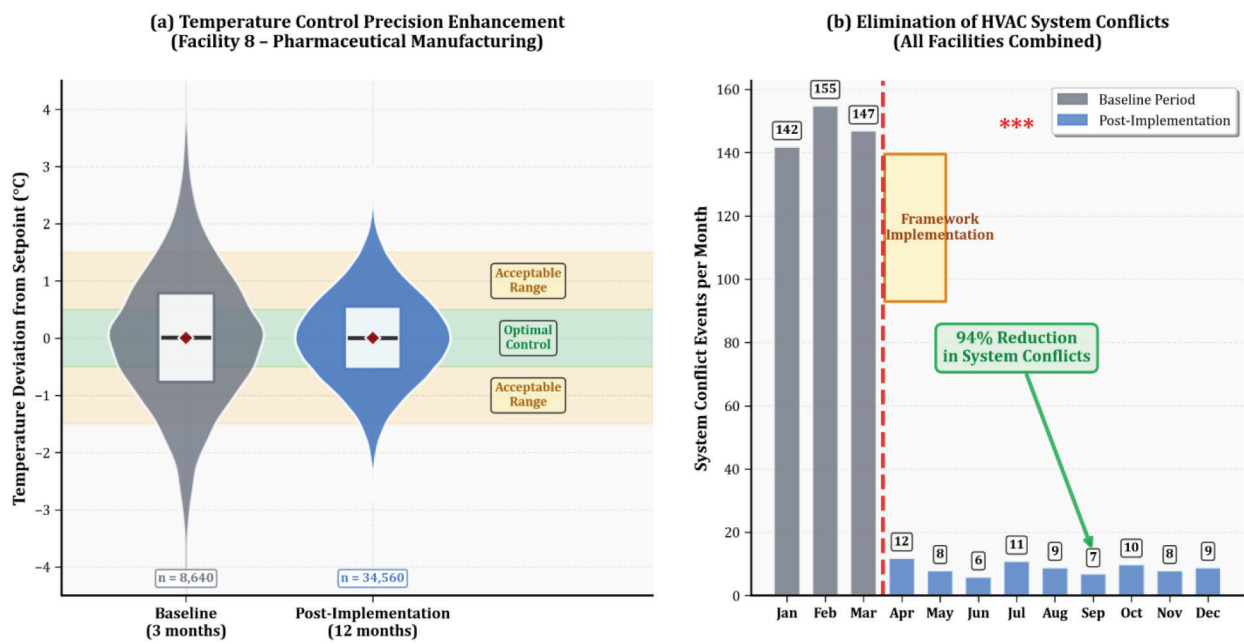


Figure 3. Improvement in Operational Control and Stability. (a) Violin plots showing the distribution of zone temperature deviations from the setpoint during baseline versus post-implementation periods, indicating improved precision. (b) Comparison of aggregate monthly system conflict events, demonstrating a 94% reduction post-implementation.

The implementation of the collaborative framework resulted in a marked improvement in the precision of environmental control. As visualized in Figure 3a, the violin plots illustrate a dramatic transformation in control stability; the broad, dispersed distribution characteristic of the baseline period (gray) is replaced by a narrow, concentrated shape in the post-implementation phase (blue). This morphological shift indicates that while baseline temperature fluctuations frequently exceeded $\pm 1.5^{\circ}\text{C}$, the collaborative control strategy successfully confined the vast majority of measurements within a tight $\pm 0.5^{\circ}\text{C}$ band. Quantitatively, the standard deviation of zone temperatures from the setpoint was reduced by an average of 37% across all facilities (from a mean of 1.25°C to 0.79°C), a statistically significant improvement ($p < 0.001$). This enhanced stability is critical in manufacturing environments like electronics and pharmaceuticals, where precise temperature and humidity control directly impacts product yield and quality.

Furthermore, the framework's ability to coordinate operations across different HVAC zones and systems virtually eliminated operational conflicts. Figure 3b shows that during the baseline period, an average of 148 system conflict events were logged per month across the eight facilities. In the baseline scenario, isolated control loops frequently engaged in "fighting," where active cooling in one high-load zone conductively lowered the temperature of an

adjacent low-load zone, inadvertently triggering the latter's heating system. This study addressed this inefficiency by incorporating inter-zonal thermal transfer coefficients into the LSTM input vector, allowing the optimizer to treat the facility as a coupled thermodynamic entity rather than disparate silos. A specific operational improvement was observed in the electronics assembly facility (FAC-01), where the model successfully predicted that heat rejection from soldering stations would naturally offset the cooling spillover from the adjacent test chamber. By anticipating this thermal balance, the framework preemptively suppressed the heater activation in the test chamber that characterized the baseline control [27]. Through this predictive coordination, the strategy reduced the occurrence of such conflicts by 94%, to an average of just 9 events per month, demonstrating a fundamental improvement in the intelligence and efficiency of the overall system logic.

4.4 Collaborative Learning Dynamics and System Evolution

A unique aspect of this study was the investigation of the FL process itself, particularly how collaboration benefited facilities with different characteristics. The analysis aimed to demonstrate that the collective intelligence of the network provided a distinct advantage over what any single facility could achieve

on its own, especially for facilities with more volatile operational profiles. The dynamics of this collaborative improvement are detailed in Figure 4. This figure's three-panel design provides a narrative of the collaborative benefit. Figure 4a tracks the monthly evolution of prediction accuracy for two facilities with disparate operational profiles: a stable chemical plant and a more volatile electronics assembly plant. Figure 4b isolates the benefit of collaboration by comparing the rate of accuracy improvement for the volatile facility under the federated model versus a simulated local-only training scenario. Finally, Figure 4c provides a conceptual visualization of the collaborative network by illustrating the relative contribution of each facility's learned insights to the final, converged global model.

The results confirm that all participants benefited from the federated approach, but the advantage was most pronounced for facilities with less predictable operational patterns. As shown in the time-series comparison in Figure 4a, the learning curves for both

the stable (Facility 4) and volatile (Facility 2) facilities exhibit a continuous upward trend in accuracy over the 12-month period. Notably, the trajectory of improvement for Facility 2 is significantly steeper, visually demonstrating its accelerated learning from the collective knowledge pool. This advantage is explicitly isolated in Figure 4b, where the red trajectory of the federated model diverges sharply from the gray baseline of the local-only simulation. This gap quantifies the "collaborative uplift," showing that the federated model for Facility 2 achieved a total accuracy improvement of 7.2 percentage points over the year, compared to only a 2.8 percentage point improvement it would have achieved training on its own data. This marked performance differential confirms that the global model successfully generalized patterns from more stable facilities to enhance predictions for the volatile loads at Facility 2.

The heatmap in Figure 4c visualizes the nature of this collaboration. It shows that while all facilities contributed to the final global model, the contributions

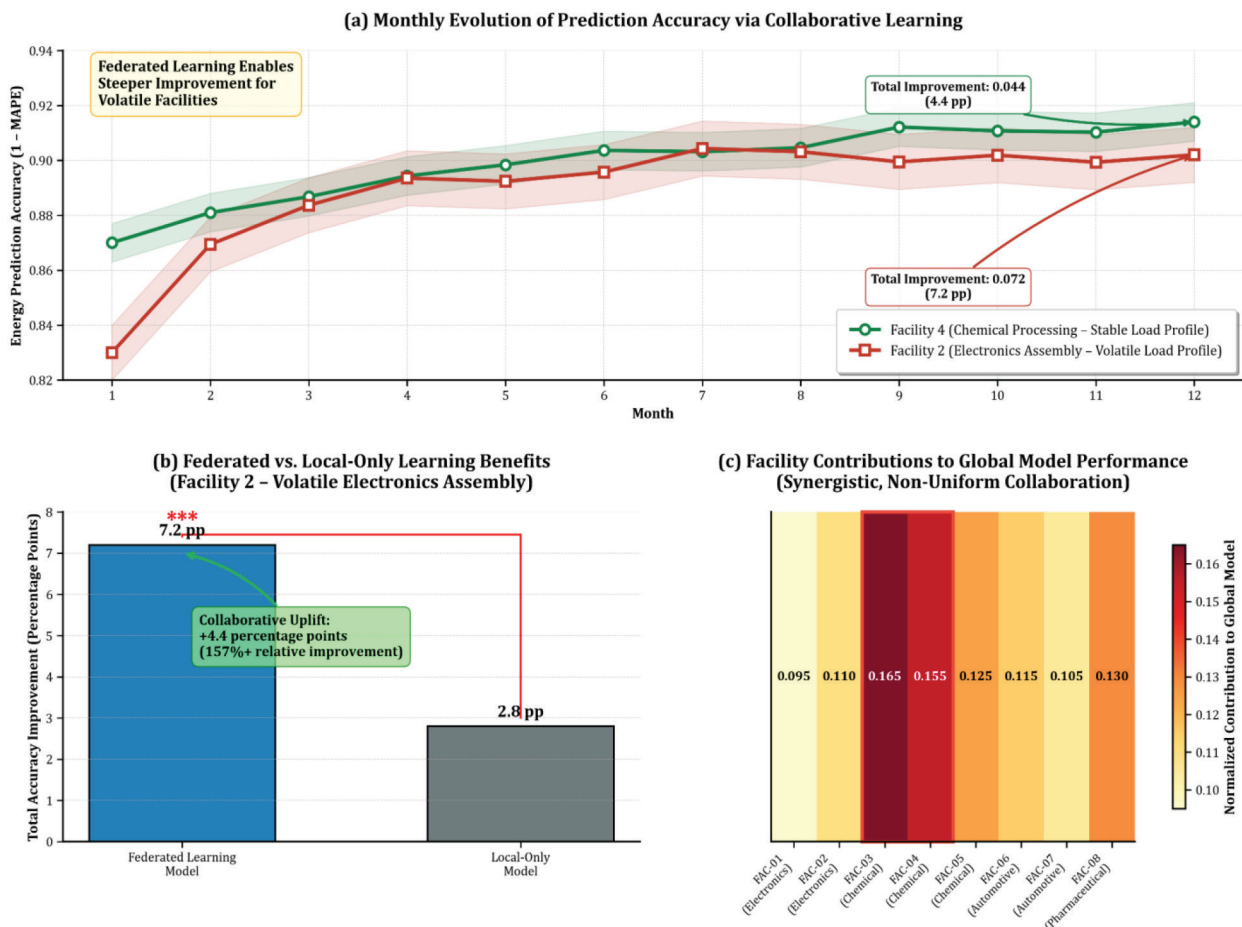


Figure 4. Dynamics of Collaborative Improvement via Federated Learning. (a) Monthly prediction accuracy improvement for facilities with stable versus volatile operational loads. (b) Comparison of total accuracy gain for a volatile-load facility under the federated model versus a local-only model. (c) Heatmap illustrating the normalized contribution of each facility to the final global model.

were not uniform. Facilities with more diverse operational data and more challenging predictive tasks (e.g., FAC-03, FAC-04) had a slightly higher impact on the final model's robustness. This indicates the creation of a truly synergistic system, where the challenges faced by one member ultimately strengthen the intelligence of the entire network [33]. Over the first year of operation, the continuous learning process enabled by the federated framework led to an aggregate system-wide performance increase of 8%, measured as an additional percentage point reduction in energy consumption from Month 1 to Month 12, on top of the initial savings.

4.5 Comparative Analysis of Framework Performance Across Facility Types

To assess the generalizability of the collaborative framework, its performance was analyzed across the four distinct manufacturing typologies represented in the study cohort. Key performance indicators were aggregated for each category to identify any systematic differences in the framework's effectiveness based on industrial application. Table 3 presents this comparative analysis, showing the mean energy savings, temperature precision improvement, and peak demand reduction for each facility type.

The results presented in Table 3 indicate that the collaborative framework delivered significant benefits across all manufacturing types, though the specific magnitude of improvements varied according to operational constraints. The highest mean energy savings (32.3%) and peak demand reductions (358.3 kW) were recorded in chemical processing facilities, where large fluid volumes provide substantial thermal storage capacity for load shifting strategies. Conversely, the pharmaceutical and electronics sectors exhibited the most pronounced improvements in temperature precision (45.0% and 41.5%, respectively), validating the system's ability to prioritize stability over raw energy minimization

in environments with stringent quality control standards. This trade-off is further reflected in the lower peak demand reduction for the pharmaceutical facility (180.0 kW), where safety margins limit the extent of pre-cooling or duty cycling permitted during peak tariff periods.

This study successfully demonstrates that a collaborative, privacy-preserving framework can significantly enhance the energy efficiency and operational stability of HVAC systems in manufacturing settings. The observed 29.1% average reduction in energy consumption and 37% improvement in temperature control precision are not merely incremental gains; they represent a paradigm shift from isolated, single-facility optimization to a more powerful system-of-systems approach. The core novelty and principal finding is that the FL architecture effectively created a collective intelligence, enabling facilities to learn from each other's operational patterns without exposing proprietary data [34]. This collaborative model's 61.8% lower prediction error compared to standalone models underscores the tangible benefits of shared learning, particularly for facilities with volatile load profiles. Table 4 provides a consolidated comparison of these aggregate performance metrics against the baseline.

These findings advance the state-of-the-art by addressing a critical gap in the literature. While previous studies have achieved notable energy savings using single-instance MPC or localized IoT analytics, they inherently operate within an information silo [17], [20], [28]. Our framework achieves comparable energy savings while adding a new dimension of inter-system coordination that prevented 94% of operational conflicts—a benefit unattainable by isolated systems. Furthermore, by employing FL, this work overcomes the significant privacy and scalability barriers that have hindered the adoption of centralized AI models in multi-firm industrial contexts, presenting a viable pathway for industry-wide collaboration [24], [46], [47].

Table 3. Comparative performance of the collaborative framework by manufacturing facility type

Facility Type	Number of Facilities	Mean Energy Savings (%)	Mean Temperature Precision Improvement (%)	Mean Peak Demand Reduction (kW)
Electronics Assembly	2	27.0%	41.5%	222.5
Chemical Processing	3	32.3%	35.0%	358.3
Automotive Parts	2	29.5%	36.5%	285.0
Pharmaceutical	1	24.0%	45.0%	180.0

Table 4. Aggregate performance comparison between baseline/local strategies and the collaborative framework

Performance Metric	Baseline / Local Control	Collaborative Framework	Net Improvement
Mean Daily Energy Consumption	61,356 kWh	43,032 kWh	29.1% Reduction
Temperature Control Precision (σ)	1.25 °C	0.79 °C	37.0% Improvement
Monthly System Conflicts (Average)	148 events	9 events	94.0% Reduction
Prediction Error (MAPE)	8.9%	3.4%	61.8% Reduction

5. Conclusions

This research successfully designed, implemented, and validated a novel collaborative framework for optimizing HVAC heat exchanger networks across multiple manufacturing facilities. This work provides conclusive evidence that a federated, predictive approach can yield substantial improvements in energy efficiency, cost-effectiveness, and operational precision, marking a significant advancement over traditional, isolated control systems. The primary empirical achievement of this work is the demonstration of significant and statistically significant energy and cost reductions. The collaborative framework achieved an average HVAC energy consumption reduction of 29.1% ($p < 0.001$) across eight diverse facilities, culminating in aggregate annualized cost savings of over \$2.1 million. This was complemented by an average peak demand reduction of 22%, showcasing the ability of the system to intelligently manage grid load. These results empirically validate the economic and operational viability of the proposed collaborative optimization strategy.

Methodologically, the principal contribution of this research is the successful application of FL to create a privacy-preserving, collective intelligence. The federated model outperformed standalone local models, achieving a 61.8% lower prediction error. This enhanced accuracy translated directly into superior operational control, evidenced by a 37% improvement in temperature control precision and a 94% reduction in wasteful system conflicts. The ability of the framework to continuously learn resulted in an additional 8% performance gain over the first year, demonstrating a robust and evolving system.

Nevertheless, the study has limitations regarding geographical scope. The cohort of eight facilities, while diverse in type, was concentrated in Saudi Arabia's Eastern Province, creating a homogenous climatic context that may limit the direct transferability

of the trained model weights to regions with disparate weather patterns. However, the underlying architecture supports adaptability through the normalization of input features; the LSTM network treats ambient conditions as dynamic variables rather than static constraints. Consequently, adaptation to contrasting climates, such as those with heating-dominant loads, would involve utilizing the global model as a pre-trained initialization point, followed by localized fine-tuning to adjust weight matrices for inverse thermal correlations. Future work should empirically validate this transfer learning capability by expanding the network to include facilities in varied climatic zones. Additionally, future research envisions integrating real-time electricity grid carbon intensity data. A proposed pilot concept involves utilizing utility APIs to ingest live emission factors, enabling the optimization algorithm to dynamically penalize consumption during high-carbon intervals. This approach effectively treats environmental impact as a distinct, weighted objective, steering the system towards not just cost efficiency but also minimized carbon footprint, thereby aligning smart manufacturing more closely with global sustainability goals. The findings confirm that collaborative IoT-enabled optimization is a powerful tool for sustainable manufacturing, effectively bridging the gap between localized control and system-wide optimization in support of Industry 5.0. While the geographical focus within Saudi Arabia represents a limitation, the architecture of the framework is inherently adaptable. Future iterations of this work will expand to varied climatic zones and integrate real-time grid carbon intensity data, steering industrial energy management toward global sustainability goals.

Funding

This research did not receive any specific grant from funding agencies in the public, commercial, or not-for-profit sectors.

References

- [1] T. Skoczkowski, S. Bielecki, M. Wołowicz, and A. Węglarz, "Redefining Energy Management for Carbon-Neutral Supply Chains in Energy-Intensive Industries: An EU Perspective," *Energies*, vol. 18, no. 15, p. 3932, 2025, doi: 10.3390/en18153932.
- [2] S. C. Vetrivel, T. P. Saravanan, and R. Maheswari, "Smart Factories and Energy Efficiency in Industry 4.0," in *Distributed Time-Sensitive Systems*, 1st ed., T. Choudhury, R. K. Singh, R. Tomar, S. Balamurugan, and J. C. Patni, Eds., Wiley, 2026, pp. 63–103. doi: 10.1002/9781139419779.ch4.
- [3] Y. Zhao, N. Li, C. Tao, Q. Chen, and M. Jiang, "A comparative study on energy performance assessment for HVAC systems in high-tech fabs," *J. Build. Eng.*, vol. 39, p. 102188, 2021, doi: 10.1016/j.jobe.2021.102188.
- [4] X. Zeng, C. Li, X. Li, C. Mao, Z. Li, and Z. Li, "Energy Efficiency Optimization of Air Conditioning Systems Towards Low-Carbon Cleanrooms: Review and Future Perspectives," *Energies*, vol. 18, no. 13, p. 3538, 2025, doi: 10.3390/en18133538.
- [5] P.-Y. Liao et al., "Energy Consumption and Carbon Emission Reduction in HVAC System of a Dynamic Random Access Memory (DRAM) Semiconductor Fabrication Plant (fab)," *IEEE Trans. Semicond. Manuf.*, vol. 37, no. 2, pp. 174–184, 2024, doi: 10.1109/TSM.2024.3379949.
- [6] N. Grasso, B. Fasciolo, A. M. M. Awouda, and G. Bruno, "A Smart Aeroponic Chamber: Structure and Architecture for an Efficient Production and Resource Management," in *Hydroponics*, N. Kumar, Ed., in *Encyclopedia of Sustainability Science and Technology Series.*, New York, NY, USA: Springer, 2024, pp. 353–380. doi: 10.1007/978-1-0716-3993-1_18.
- [7] A. Alrabghi, "A modelling approach for asset degradation: Advancing digital twin in maintenance," *Int. J. Simul. Model.*, vol. 24, no. 1, pp. 76–86, Mar. 2025, doi: 10.2507/IJSIMM24-1-715.
- [8] M. A. Alghassab, "Fuzzy-based smart energy management system for residential buildings in Saudi Arabia: A comparative study," *Energy Rep.*, vol. 11, pp. 1212–1224, 2024, doi: 10.1016/j.egy.2023.12.039.
- [9] T. M. Olatunde, A. C. Okwandu, D. O. Akande, and Z. Q. Sikhakhane, "Review of energy-efficient HVAC technologies for sustainable buildings," *Int. J. Sci. Technol. Res. Arch.*, vol. 6, no. 2, pp. 12–20, 2024, doi: 10.53771/ijstra.2024.6.2.0039.
- [10] A. Madadzadeh, K. Siddiqui, and A. A. Aliabadi, "The economics landscape for building decarbonization," *Sustainability*, vol. 16, no. 14, p. 6214, 2024, doi: 10.3390/su16146214.
- [11] I. Spasojevic, S. Havzi, D. Stefanovic, S. Ristic, and U. Marjanovic, "Research Trends and Topics in IJIEM from 2010 to 2020: A Statistical History," *Int. J. Ind. Eng. Manag.*, vol. 12, no. 4, pp. 228–242, 2021, doi: 10.24867/IJIEM-2021-4-290.
- [12] Y. Gao, S. Li, Y. Xiao, W. Dong, M. Fairbank, and B. Lu, "An iterative optimization and learning-based IoT system for energy management of connected buildings," *IEEE Internet Things J.*, vol. 9, no. 21, pp. 21246–21259, 2022, doi: 10.1109/JIOT.2022.3176306.
- [13] M. Siahkouhi, M. Rashidi, F. Mashiri, F. Aslani, and M. S. Ayubirad, "Application of self-sensing concrete sensors for bridge monitoring- A review of recent developments, challenges, and future prospects," *Measurement*, vol. 245, p. 116543, 2025, doi: 10.1016/j.measurement.2024.116543.
- [14] A. Yayla et al., "Artificial intelligence (AI)-based occupant-centric heating ventilation and air conditioning (HVAC) control system for multi-zone commercial buildings," *Sustainability*, vol. 14, no. 23, p. 16107, 2022, doi: 10.3390/su142316107.
- [15] D. K. Priatna, W. Roswinna, N. Limakrisna, A. Khalikov, D. Abdullaev, and L. Hussein, "Optimizing Smart Manufacturing Processes and Human Resource Management through Machine Learning Algorithms," *Int. J. Ind. Eng. Manag.*, vol. 16, no. 2, pp. 176–188, 2025, doi: 10.24867/IJIEM-382.
- [16] G. Nota, F. D. Nota, A. Toro, and M. Nastasia, "A framework for unsupervised learning and predictive maintenance in Industry 4.0," *Int. J. Ind. Eng. Manag.*, vol. 15, no. 4, pp. 304–319, 2024, doi: 10.24867/IJIEM-2024-4-365.
- [17] F. Ascione, R. F. De Masi, V. Festa, G. M. Mauro, and G. P. Vanoli, "Optimizing space cooling of a nearly zero energy building via model predictive control: Energy cost vs comfort," *Energy Build.*, vol. 278, p. 112664, 2023, doi: 10.1016/j.enbuild.2022.112664.
- [18] A. G. Mainini et al., "Enabling Sufficiency Through Smart Locks: Transforming Office Occupancy and Building Management for Energy Savings," *Buildings*, vol. 15, no. 5, p. 669, 2025, doi: 10.3390/buildings15050669.
- [19] D. Huber, "Multi-objective heat exchanger network synthesis," Ph.D. dissertation, Technische Universität Wien, Vienna, Austria, 2024. [Online]. Available: <https://repositum.tuwien.at/handle/20.500.12708/198925>.
- [20] J. C. Atuonwu, "A simulation tool for pinch analysis and heat exchanger/heat pump integration in industrial processes: Development and application in challenge-based learning," *Educ. Chem. Eng.*, vol. 52, pp. 141–150, 2025, doi: 10.1016/j.ece.2025.04.001.
- [21] H. A. Er, S. R. Wan Alwi, and Z. A. Manan, "A graphical method for Heat Exchanger Storage Network synthesis to decarbonise the non-continuous industrial processes," *Clean Technol. Environ. Policy*, vol. 27, pp. 8599–8622, 2024, doi: 10.1007/s10098-024-02983-z.
- [22] S. H. Godasiaei, "Exploring the influence of crystallization fouling on microscale heat exchangers through machine learning analysis," *Numer. Heat Transf. A*, vol. 86, no. 22, pp. 8022–8048, 2024, doi: 10.1080/10407782.2024.2357582.
- [23] R. Sukarno, N. Putra, I. I. Hakim, F. F. Rachman, and T. M. I. Mahlia, "Utilizing heat pipe heat exchanger to reduce the energy consumption of airborne infection isolation hospital room HVAC system," *J. Build. Eng.*, vol. 35, p. 102116, 2021, doi: 10.1016/j.jobe.2020.102116.
- [24] L. Albshaiyer, S. Almarri, and A. Albulai, "Federated learning for cloud and edge security: A systematic review of challenges and AI opportunities," *Electronics*, vol. 14, no. 5, p. 1019, 2025, doi: 10.3390/electronics14051019.
- [25] F. Hagström, V. Garg, and F. Oliveira, "Employing federated learning for training autonomous HVAC systems," *Energy Build.*, vol. 340, p. 115761, 2025, doi: 10.1016/j.enbuild.2025.115761.
- [26] G. Feng, "Optimization Control of HVAC System and Building Energy Management Based on Machine Learning," in *Proc. IEEE CSNT*, 2025, pp. 263–268, doi: 10.1109/CSNT64827.2025.10967665.
- [27] N. S. Raman, R. U. Chaturvedi, Z. Guo, and P. Barooah, "Model predictive control-based hierarchical control of a multi-zone commercial HVAC system," *J. Eng. Sustain. Build. Cities*, vol. 2, no. 2, p. 021005, 2021, doi: 10.1115/1.4051205.
- [28] H. Tan et al., "Combining reinforcement learning with mathematical programming: An approach for optimal design of heat exchanger networks," *Chin. J. Chem. Eng.*, vol. 69, pp. 63–71, 2024, doi: 10.1016/j.cjche.2023.12.005.

- [29] G. Lee, Y. Joo, S.-U. Lee, T. Kim, Y. Yu, and H.-G. Kim, "Design optimization of heat exchanger using deep reinforcement learning," *Int. Commun. Heat Mass Transf.*, vol. 159, p. 107991, 2024, doi: 10.1016/j.icheatmasstransfer.2024.107991.
- [30] H. K. Channi and C. L. Chowdhary, "Deep Learning Approach Towards Green IIOT," in *Smart Computing Techniques in Industrial IoT*, vol. 1172, C. L. Chowdhary, A. K. Tripathy, and Y. Wu, Eds., in *Studies in Computational Intelligence*, vol. 1172, Singapore: Springer Nature Singapore, 2025, pp. 115–142. doi: 10.1007/978-981-97-7494-4_7.
- [31] M. A. Khan et al., "Smart buildings: Federated learning-driven secure, transparent and smart energy management system using XAI," *Energy Rep.*, vol. 13, pp. 2066–2081, 2025, doi: 10.1016/j.egy.2025.01.063.
- [32] B. Amangeldy, T. Imankulov, N. Tasmurzayev, G. Dikhanbayeva, and Y. Nurakhov, "A Review of Artificial Intelligence and Deep Learning Approaches for Resource Management in Smart Buildings," *Buildings*, vol. 15, no. 15, p. 2631, 2025, doi: 10.3390/buildings15152631.
- [33] Cheng, X., Li, C., and Liu, X., "A review of federated learning in energy systems," in *Proc. IEEE/IAS I&CPS Asia*, Shanghai, China, 2022, pp. 2089–2095, doi: 10.48550/arXiv.2208.10941.
- [34] T. Wang, X. Zuo, P. Wang, Z. Zhang, and D. Wang, "A Data Optimization and Updating Method for Single Building," *Sens. Mater.*, vol. 37, no. 1(3), pp. 231–245, 2025, doi: 10.18494/SAM5290.
- [35] J. Akbari and M. S. Ayubirad, "Seismic Optimum Design of Steel Structures Using Gradient-Based and Genetic Algorithm Methods," *Int. J. Civ. Eng.*, vol. 15, no. 2, pp. 135–148, 2017, doi: 10.1007/s40999-016-0088-0.
- [36] I. Ficili, M. Giacobbe, G. Tricomi, and A. Puliafito, "From Sensors to Data Intelligence: Leveraging IoT, Cloud, and Edge Computing with AI," *Sensors*, vol. 25, no. 6, p. 1763, 2025, doi: 10.3390/s25061763.
- [37] U. Kubayev et al., "Adaptive Islanding Detection in Microgrids Using Deep Learning and Fuzzy Logic for Enhanced Stability and Accuracy," *J. Oper. Autom. Power Eng.*, vol. 12, no. Special Issue, pp. 33–42, 2024, doi: 10.22098/joape.2025.16153.2247.
- [38] C. Zhang, Y. Xie, H. Bai, B. Yu, W. Li, and Y. Gao, "A survey on federated learning," *Knowl.-Based Syst.*, vol. 216, p. 106775, 2021, doi: 10.1016/j.knosys.2021.106775.
- [39] K. Bonawitz, P. Kairouz, B. McMahan, and D. Ramage, "Federated Learning and Privacy: Building privacy-preserving systems for machine learning and data science on decentralized data," *Queue*, vol. 19, no. 5, pp. 87–114, 2021, doi: 10.1145/3494834.3500240.
- [40] S. Bharati, M. R. H. Mondal, P. Podder, and V. B. S. Prasath, "Federated learning: Applications, challenges and future directions," *Int. J. Hybrid Intell. Syst.*, vol. 18, no. 1–2, pp. 19–35, 2022, doi: 10.3233/HIS-220006.
- [41] E. T. M. Beltrán et al., "Decentralized federated learning: Fundamentals, state of the art, frameworks, trends, and challenges," *IEEE Commun. Surv. Tutor.*, vol. 25, no. 4, pp. 2983–3013, 2023, doi: 10.1109/COMST.2023.3315746.
- [42] D. Zhuang, V. J. Gan, Z. D. Tekler, A. Chong, S. Tian, and X. Shi, "Data-driven predictive control for smart HVAC system in IoT-integrated buildings with time-series forecasting and reinforcement learning," *Appl. Energy*, vol. 338, p. 120936, 2023, doi: 10.1016/j.apenergy.2023.120936.
- [43] H. Alizadegan, B. Rashidi Malki, A. Radmehr, H. Karimi, and M. A. Ilani, "Comparative study of long short-term memory (LSTM), bidirectional LSTM, and traditional machine learning approaches for energy consumption prediction," *Energy Explor. Exploit.*, vol. 43, no. 1, pp. 281–301, 2025, doi: 10.1177/01445987241269496.
- [44] S. Algburi et al., "Predictive modeling of building energy consumption and thermal comfort for decarbonization in construction and retrofitting," *Results Eng.*, vol. 26, p. 105475, 2025, doi: 10.1016/j.rineng.2025.105475.
- [45] Z. M. S. El-Barbary, L. Safarova, F. Atamurotov, A. M. Alsayah, and B. K. Yadav, "Artificial Intelligence Driven Internet of Things Framework for Wind Energy Monitoring and Performance Enhancement in Smart Cities," *Int. J. Energy Res.*, vol. 2025, no. 1, p. 7215655, 2025, doi: 10.1155/er/7215655.
- [46] S. Algburi et al., "The role of artificial intelligence in accelerating renewable energy adoption for global energy transformation," *Unconv. Resour.*, vol. 8, p. 100229, 2025, doi: 10.1016/j.unres.2025.100229.
- [47] J. Begatov, A. Mukhamedov, M. Platoshina, S. Saidova, and M. Yakubova, "Formation of The Structure of Tool Steels During Heat Treatment with Heating to Extreme Temperatures," *Int. J. Mechatron. Appl. Mech.*, vol. 17, pp. 153–158, 2024.