Megaron
https://megaron.yildiz.edu.tr - https://megaronjournal.com
DOI: https://doi.org/10.14744/megaron.2025.02256

MEGARON

Article

Supervised machine learning for thermal comfort and energy efficiency: An evaluation for the indoor built environment

Ali Berkay AVCI*

Department of Architecture, Süleyman Demirel University, Isparta, Türkiye

ARTICLE INFO

Article history
Received: 06 April 2025
Revised: 09 October 2025
Accepted: 09 October 2025

Key words:
Energy efficiency; HVAC; machine learning; neural networks; supervised learning; thermal comfort.

ABSTRACT

The growing demand for energy-efficient and sustainable buildings has accelerated the exploration of advanced technologies to optimize thermal comfort and reduce energy consumption. Machine learning techniques, particularly supervised learning approaches, have shown strong potential to optimize HVAC control while maintaining comfort. However, existing studies are often fragmented, with limited integrated analyses of methodologies and applications, particularly in the context of diverse climates, building typologies, and occupant behaviors. This study addresses these gaps through a semi-systematic review of peer-reviewed studies applying supervised machine learning techniques for thermal comfort prediction and energy optimization. Using a transparent process involving Web of Science search, predefined inclusion/exclusion criteria, and Rayyan-assisted screening, 18 supervised learning articles were identified from an initial 603 records. These articles were categorized into tree-based models, regression-based models and neural networks. The review identifies critical gaps, such as the insufficient integration of real-time occupant behavior, limited applicability across diverse climatic conditions, and challenges in achieving a balance between energy efficiency and occupant comfort. Findings highlight the strengths of tree-based models in feature selection and real-time decision-making, the simplicity of regression-based models for controlled environments, and the adaptability of neural networks in complex, non-linear scenarios. Despite these advancements, limitations such as data scarcity, computational demands, and the lack of long-term validation persist. Addressing these challenges is essential for the development of robust and scalable machine learning-driven solutions. This study provides a roadmap for future research and practical applications, emphasizing the transformative potential of supervised machine learning techniques in achieving sustainable, energy-efficient, and occupant-centered building environments.

Cite this article as: Avcı A.B. (2025). Supervised machine learning for thermal comfort and energy efficiency: An evaluation for the indoor built environment. Megaron, 20(3), 418-432.

INTRODUCTION

The increasing global demand for energy-efficient and sustainable buildings has driven the exploration of advanced technologies for optimizing thermal comfort and reducing energy consumption (Moshood et al., 2024). As heating, ventilation, and air-conditioning (HVAC) systems account for a significant proportion of energy use in buildings (Gupta & Deb, 2022), there is a critical need to develop intelligent control systems capable of maintaining optimal indoor environments while minimizing energy costs (Halhoul Merabet et al., 2021). Supervised machine

***Corresponding author**
*E-mail adres: aliberkayavci@gmail.com

Published by Yıldız Technical University, İstanbul, Türkiye

learning techniques have emerged as powerful tools in this domain (Essamlali et al., 2024). They provide data-driven approaches to predict and manage indoor thermal conditions. These methods enable dynamic adjustments in HVAC operations, ensuring a balance between energy efficiency and occupant comfort across diverse building typologies and climatic conditions (Zhou et al., 2023).

This study focuses on three supervised machine learning (ML) families that have demonstrated substantial promise in addressing challenges related to thermal comfort control and energy optimization. Researchers aim to enhance the adaptability and precision of HVAC systems by leveraging these models. These attributes make them responsive to both environmental variations and occupant behavior. Despite their advancements, the application of supervised machine learning techniques remains an evolving field, requiring a comprehensive understanding of their strengths, limitations, and practical implications. To address this need, a semi-systematic literature review was undertaken using explicit eligibility criteria, a Web of Science database search, Rayyan-assisted screening, and structured data extraction.

Aim of the study

The exploration of machine learning in thermal comfort prediction has produced a wide range of studies, but the fragmented nature of existing research leaves significant opportunities for synthesis and further innovation. Several previous studies have focused on reviewing thermal comfort prediction studies using machine learning algorithms. Qavidel Fard et al. (2022) conducted a systematic review focusing on ML applications in thermal comfort studies, emphasizing methods, performance, and challenges. However, their review identified a lack of focus on personal comfort models and inadequate exploration of real-world applications, alongside challenges in feature selection and model evaluation. Similarly, Feng et al. (2022) reviewed data-driven methods for personal thermal comfort prediction, addressing experimental design and modeling techniques but found insufficient attention to inter- and intra-individual variability and limited integration of online learning techniques. Another review by Lala & Hagishima (2022) provided a unique perspective on thermal comfort for primary schools, focusing on children-specific challenges such as illogical votes, multiple comfort metrics, and class imbalance in machine learning prediction studies. Yet, they emphasized the absence of dedicated machine learning models for children, indicating a gap in integrating these insights into broader contexts. Arakawa Martins et al. (2022) systematically reviewed personal thermal comfort models but identified limited diversity in climatic conditions, building typologies, and participant demographics, coupled with challenges in standardizing predictive frameworks. Finally, Zhang et al. (2022) critically reviewed machine learning-based occupancy prediction models, linking energy efficiency and indoor environmental quality. However, their analysis highlighted gaps in integrating occupancy prediction with real-time HVAC optimization and the need for addressing perceived indoor air quality (IEQ) and thermal comfort jointly.

These reviews' findings emphasize the necessity for a review that synthesizes insights from supervised machine learning techniques in predicting indoor thermal comfort while addressing their limitations and bridging the identified gaps. Table 1 summarizes key review studies, highlighting their purposes and the research gaps identified. The present research aims to evaluate the effectiveness of supervised

Table 1. Summary of the recent review papers on machine learning in thermal comfort

Review Study	Purpose of the Review	Identified Gaps
(Qavidel Fard et al., 2022)	Systematic review of ML applications in thermal comfort to evaluate methods, performance, and challenges.	Lack of focus on personal comfort models, inadequate exploration of real-world applications, and challenges in feature selection and model evaluation.
(Feng et al., 2022)	Review of data-driven methods for personal thermal comfort prediction, focusing on experimental design, data collection, and modeling techniques.	Insufficient attention to inter- and intra-individual variability, data quality issues, and limited integration of online learning techniques.
(Lala & Hagishima, 2022)	Comprehensive review of thermal comfort in primary schools, addressing ML challenges specific to children.	Absence of dedicated ML models for children, challenges like illogical votes, and data imbalance in primary school thermal comfort prediction.
(Arakawa Martins et al., 2022)	Systematic review of personal thermal comfort models with a focus on predictive modeling processes.	Limited diversity in climate, building types, and participant demographics; challenges in standardizing predictive modeling frameworks.
(Zhang et al., 2022)	Critical review of ML-based occupancy prediction models for energy efficiency, air quality, and thermal comfort.	Gaps in integrating occupancy prediction with real-time HVAC optimization and limited studies addressing perceived IEQ and thermal comfort jointly.

machine learning techniques in optimizing energy use and maintaining thermal comfort within built environments. By categorizing and analyzing research articles into three ML approaches, the study explores methodologies, applications, and outcomes, seeking to:

- Elucidate the contributions of each machine learning approach to thermal comfort optimization.
- Identify and address limitations and inconsistencies in existing research.
- Provide actionable insights to advance sustainable building practices through machine learning-driven HVAC systems.

METHODOLOGY

This study adopted a semi-systematic review approach to balance transparency and reproducibility with a focused scope on supervised machine learning for indoor thermal comfort and energy optimization. The review followed a series of predefined steps to ensure the selection of high-quality and relevant literature, detailed as follows:

- **Data source and search strategy:** Articles indexed in the Web of Science Core Collection were selected as the data source due to its comprehensive indexing of peer-reviewed scientific publications, ensuring access to high-impact studies. A database search was performed using the keywords "Thermal Comfort" and "Machine Learning," yielding 603 articles, without applying any year limit. This focused selection ensured methodological consistency and avoided redundancy across overlapping indexing platforms, which often contain identical records within this specialized research area. The choice of a single, high-quality source also enabled a transparent and reproducible workflow, emphasizing depth and reliability over breadth of coverage.
- **Inclusion and Exclusion Criteria:** Several inclusion and exclusion criteria were applied to refine the dataset. Inclusion criteria were peer-reviewed journal articles addressing indoor built environments, supervised machine learning models, and outcomes on indoor thermal comfort and/or energy demand. The initial exclusion criteria were review papers, non-English or non-open-access items, theses/abstracts/grey literature to maintain the focus on indoor built environments. The exclusion criteria for the first screening process were personal comfort systems (PCSs) relying on physiological data, outdoor thermal comfort (OTC), vehicle indoor environments, and other non-indoor built environment research. Then, for the second screening session, the studies using unsupervised learning, reinforcement learning, and hybrid ML methods were excluded. These criteria were established to ensure reproducibility and enable further research.
- **Screening:** From 603 records, total 354 records were excluded, which were 60 review studies, 114 proceeding papers, three non-English studies, and 177 non-open access papers. The remaining 249 records were imported into Rayyan, a collaborative systematic review tool, to enhance the efficiency of the screening process (Ouzzani et al., 2016). Abstracts of 249 records were screened based on the aforementioned inclusion and exclusion criteria. At this stage, 215 articles were excluded because they did not address supervised learning for indoor built environments. These comprised 93 studies on personal comfort systems (PCSs) relying on physiological data, 59 on outdoor thermal comfort (OTC), 36 on vehicle indoor environments, and 27 on other non-indoor built environment contexts. This refinement resulted in 34 research articles focusing on machine learning techniques applied to thermal comfort models and energy efficiency. A second round of screening was then conducted on the full texts of these 34 articles to retain only those employing supervised learning techniques. In this stage, 16 articles were excluded: Eight focused on unsupervised learning methods, five on reinforcement learning methods, and three on hybrid machine learning approaches. The second screening produced 18 included studies specifically employing supervised learning techniques for indoor thermal comfort (603 → 249 → 34 → 18). The decision to focus on supervised learning was based on its dominance and preference in the field (Han et al., 2023), given its ability to handle labeled data for predictive accuracy and its wide applicability in real-world HVAC systems. Other machine learning approaches, including unsupervised and reinforcement learning, while valuable, were less represented and often lacked the direct applicability to thermal comfort optimization within building environments (Zhang et al., 2022).
- **Data extraction and categorization:** After screening, the included studies were first categorized according to the type of supervised machine learning approach used to facilitate comparative analysis:
  - **Tree-Based Models:** These models, including Random Forest and Gradient Boosted Decision Trees, were evaluated for their interpretability and robustness in handling diverse datasets.
  - **Regression-Based Models:** Studies focusing on linear and non-linear regression techniques were analyzed for their simplicity and adaptability in predicting thermal comfort indices.
  - **Neural Network Applications:** Advanced neural network architectures, including Artificial Neural Networks (ANNs) and Physics-Informed Neural Networks (PiNNs), were reviewed for their ability to model complex, non-linear relationships.

Following this categorization, structured data extraction was undertaken for each included study in each category to ensure consistency and comparability. Predefined fields included context/building type, climate/geographical setting, dataset size and variables (environmental and occupant-related parameters), validation method (cross-validation, field testing, simulation-based), performance metrics (accuracy, F1-score, MAE, RMSE, energy savings), and energy/comfort outcomes (PMV, TSV, operative temperature, optimization levels, comfort improvements).

- **Quality Assessment / Risk of Bias:** To evaluate the methodological quality and reliability of the included studies, the Prediction model Risk of Bias Assessment Tool (PROBAST) was adapted to the context of supervised machine learning in thermal comfort and energy optimization (Wolff et al., 2019). Each study was independently assessed across four domains, data and setting (D1), predictors and feature engineering (D2), outcomes/labels (D3), and analysis (D4), and rated as low, low–moderate, moderate, or high risk of bias. The applicability concerns were also rated and recorded a short justification for each judgment (Appendix Table A1). This approach ensured transparency and reproducibility in the evaluation of the included studies.
- **Synthesis and comparative analysis:** A structured narrative synthesis was undertaken, organizing the included studies according to the three supervised machine learning approach groups and enabling systematic cross-study comparison of methods and outcomes. This process allowed the identification of patterns, methodological differences, and performance trends across the reviewed studies. The main analytical dimensions included:
  - **Adaptability:** Assessing the ability of models to adjust to varying climates, building typologies, and occupant behaviors in real-time settings.
  - **Challenges:** Identifying specific limitations such as insufficient real-time data processing or incomplete integration of occupant behavior.
  - **Climate influence:** Examining how differences in climatic conditions affected optimization levels, highlighting areas where machine learning models underperform.

This synthesis and comparative analysis addressed critical questions such as whether the limitations stemmed from inadequate real-time data integration or inherent gaps in capturing occupant behavior. Thus, this approach provides a deeper understanding of the research landscape and identifies avenues for further development.

The structured selection process for the reviewed studies is illustrated in Figure 1, providing a visual representation of each searching stage, screening, and eligibility assessment, and categorization phase.

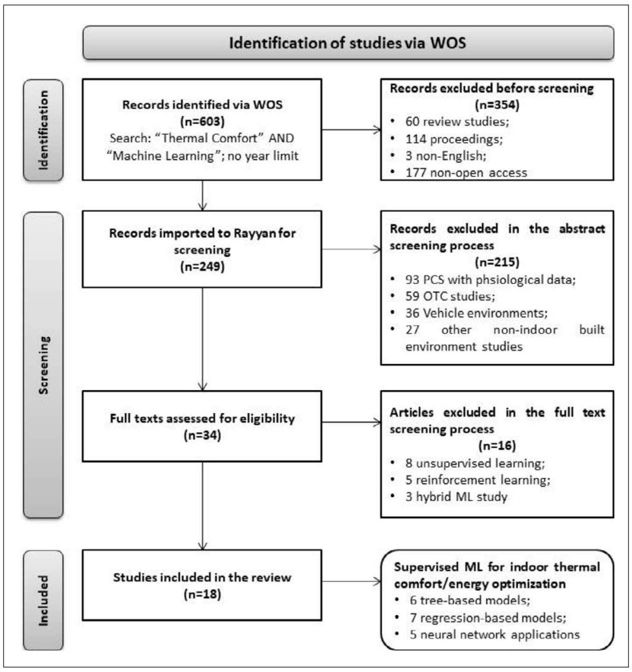

Figure 1. Flow chart of the selection strategy and categorization process.

RESULTS AND DISCUSSION

Categorization of supervised learning approaches

Supervised learning techniques have been widely used in the field of thermal comfort prediction and energy optimization within the built environment. These methods employ labeled data to establish predictive models that address diverse challenges such as real-time HVAC control, adaptive thermal comfort management, and efficient energy use. The reviewed studies are categorized into three primary approaches: Tree-based models, regression-based models, and neural network applications. Each category demonstrates distinct strengths and applications, from robust feature selection in tree-based models to the simplicity and interpretability of regression-based approaches, and the powerful adaptability of neural networks in handling complex, non-linear datasets. This section provides a detailed analysis of the selected articles under these categories, exploring their methodologies, applications, and outcomes, while also highlighting their contributions to addressing challenges in thermal comfort control and energy efficiency.

Tree-Based Models: Tree-based machine learning models have garnered attention for their ability to process complex datasets efficiently, offering robust feature selection and interpretability. This section details the methodologies and findings of six key studies that utilized tree-based approaches for thermal comfort and energy optimization. These studies denote the versatility of models such as Random Forest (RF), Gradient Boosted Decision Trees

(GBDT), and Decision Trees (DT) across different thermal comfort applications. Aparicio-Ruiz et al. focused on understanding indoor thermal comfort in Mediterranean climates using RF. This study emphasized the importance of an extended database with 21 variables, including indoor and outdoor parameters such as CO2 levels and running mean temperature. By employing RF, the researchers achieved a 5% improvement in accuracy, illustrating the RF's capacity to handle diverse inputs and optimize conditioning systems for Mediterranean buildings (Aparicio-Ruiz et al., 2023). Similarly, Bai et al. compared the performance of RF and GBDT in predicting thermal preferences based on the ASHRAE Comfort Database II. Their ensemble learning approach demonstrated robust results, achieving weighted F1-scores more than 90%. The study also highlighted the influence of data characteristics like building type and season on model performance, showcasing the adaptability of tree-based models in varying contexts (Bai et al., 2022). On the other hand, Brik et al. integrated RF with Internet of Things (IoT) frameworks to create a real-time thermal comfort monitoring system. They achieved a prediction accuracy of 96% and reduced deviation from setpoints by 85% using data from a longitudinal study. Their study illustrated the synergy between IoT technologies and machine learning, offering insights into energy-efficient building management (Brik et al., 2022).

Hosamo et al. introduced an innovative application of RF within a digital twin framework to improve predictive maintenance and occupant comfort. The integration of Building Information Modelling (BIM) and real-time sensor data allowed for accurate detection of HVAC issues, reducing system failures by 10%. This study underlined the potential of RF in advancing maintenance strategies and enhancing energy efficiency (Hosamo et al., 2023). In the study of Lu et al. RF model was applied to a combined radiant floor and fan coil cooling system, focusing on predicting operative temperature and energy consumption. Their findings demonstrated that RF outperformed other algorithms in error metrics, with reductions of up to 82% in mean squared error. The study emphasized the importance of machine learning in optimizing hybrid cooling systems, particularly in high-demand scenarios (Lu et al., 2024). Finally, Miao et al. developed an RF-based model tailored for naturally ventilated educational buildings. The study identified occupancy and ventilation practices as critical factors influencing thermal comfort. By leveraging accessible data, the researchers provided a cost-effective solution for schools, achieving robust generalization and practical applicability without the need for extensive sensor networks (Miao et al., 2023). These studies demonstrate the versatility and efficacy of tree-based models in addressing challenges related to thermal comfort and energy optimization. The details of the methodologies and findings are summarized in Table 2.

Regression-Based Models: Regression-based models serve as fundamental tools in predicting thermal comfort metrics by combining simplicity and interpretability. These models stand out in analyzing the relationships between environmental and personal factors with thermal comfort indices like Predicted Mean Vote (PMV) and Thermal Sensation Vote (TSV). This section synthesizes seven determined studies, their methodologies, and findings, elaborating on their contributions to this field. Abdellatif et al. presented a hybrid methodology integrating Multiple

Table 2. Summary of reviewed articles employing tree-based models

Study	Model	Application	Key Metrics	Outcomes
(Aparicio-Ruiz et al., 2023)	Random Forest	Thermal comfort prediction in Mediterranean climates	Accuracy Improvement: 5%	Enhanced model performance and variable relevance identification.
(Bai et al., 2022)	RF, GBDT	Thermal preference prediction	F1-Score: >90%	Superior performance of ensemble models with expanded datasets.
(Brik et al., 2022)	Random Forest	Real-time IoT-based thermal comfort monitoring	Accuracy: 96%, Optimization: 85%	Improved indoor comfort and real-time optimization capabilities.
(Hosamo et al., 2023)	Random Forest	Digital twin-based predictive maintenance	HVAC Failures Reduced: 10%	Enhanced occupant comfort and equipment lifespan through predictive strategies.
(Lu et al., 2024)	Random Forest	Hybrid cooling system energy and comfort prediction	MSE Improvement: 82%	Significant energy savings and predictive accuracy in hybrid cooling systems.
(Miao et al., 2023)	Random Forest	Educational building thermal comfort	Robust Generalization	Cost-effective prediction models for schools relying on natural ventilation.

Linear Regression (MLR) with genetic algorithms to optimize heating strategies for office buildings. By employing data from TRNSYS simulations, their approach achieved a 43% improvement in thermal comfort while maintaining energy efficiency. The genetic algorithm optimized heating parameters, demonstrating the utility of regression models in dynamic control systems (Abdellatif et al., 2022). Kumar & Kurian (2023) explored real-time thermal comfort prediction using Bayesian-optimized regression models. Their study developed predictive tools for PMV and Standard Effective Temperature (SET), leveraging automated feature selection techniques like Neighborhood Component Analysis. This model enhanced HVAC system responsiveness, yielding significant energy savings and improved user satisfaction through real-time environmental adjustments. Another study by Liu & Ma (2023) proposed an explainable Light Gradient Boosted Machine (LightGBM) regression model combined with SHAP analysis to assess thermal comfort across diverse Chinese climates. Their approach provided interpretable insights into the interactive effects of building and climatic variables, facilitating region-specific design optimizations aligned with energy conservation goal.

Mousavi et al. (2023) utilized meta-additive regression within a Green Building framework to optimize residential building envelopes in semi-arid climates. This study employed DesignBuilder simulations and statistical optimization to determine the most effective combinations of envelope parameters. Their methodology led to substantial annual energy reductions, emphasizing the adaptability of regression models in passive design strategies. Park et al. (2024) conducted a field test integrating MLR within a thermal comfort controller (TCC) for residential HVAC systems. Their model utilized mean radiant temperature estimations to achieve real-time adjustments in HVAC settings, resulting in a 60% improvement in PMV and over 20% energy savings. This study highlighted the effectiveness of regression in real-world applications under dynamic climatic conditions. Similarly, Sibyan et al. (2022) compared MLR with machine learning approaches like Naive Bayes classifiers for thermal comfort prediction in naturally ventilated environments. The analysis demonstrated MLR's robustness despite simpler assumptions, validating its applicability in field studies and comparative analyses. Finally, Xi et al. (2024) applied linear regression to assess TSV in traditional Chinese dwellings. This study integrated field measurements and subjective surveys, identifying temperature ranges that aligned with historical and modern thermal comfort requirements. Their findings underscored the importance of contextual factors, such as cultural preferences and architectural heritage, in predictive modeling. These studies collectively underline the versatility of regression-based models in addressing thermal comfort challenges across various contexts. The methodologies and outcomes of these studies are detailed in Table 3.

Neural Network Applications: Artificial Neural Networks (ANNs) have emerged as a pivotal tool in advancing thermal comfort prediction and energy optimization within building management systems. By effectively modelling non-linear and complex relationships among environmental and personal parameters, ANNs demonstrate significant advantages in handling diverse datasets and achieving high predictive accuracy. In this section, the methodologies and

Table 3. Summary of reviewed articles employing regression-based models

Study	Regression Model	Application	Key Metrics	Outcomes
(Abdellatif et al., 2022)	Multiple Linear Regression	Heating optimization for indoor comfort	<1% Error, Adjusted R2: 0.9	43% improvement in thermal comfort, significant energy savings.
(Kumar & Kurian, 2023)	Bayesian-Optimized MLR	Real-time PMV and SET prediction	High Accuracy, Fast Response	Enhanced HVAC efficiency, real-time adaptability.
(Liu & Ma, 2023)	LightGBM Regression	Thermal comfort evaluation across climates	Accuracy with SHAP interpretations	Improved thermal designs for regional diversity.
(Mousavi et al., 2023)	Meta-Additive Regression	Envelope optimization in semi-arid climates	50% Energy Reduction	Optimal building design for energy and comfort enhancement.
(Park et al., 2024)	Linear Regression	Real-time HVAC control	PMV: +60%, Energy Savings: >20%	Improved comfort and efficiency in hot-dry climates.
(Sibyan et al., 2022)	Multiple Linear Regression	Comparison with ML methods	Higher Prediction Accuracy	Validated regression accuracy in field studies.
(Xi et al., 2024)	Linear Regression	TSV prediction for heritage dwellings	Accurate TSV Models	Novel insights for heritage building thermal comfort.

findings of five key studies that illustrate the application of ANNs are presented. Boutahri & Tilioua highlighted the predictive capabilities of ANNs in forecasting PMV values with enhanced accuracy, achieving an energy-saving potential of up to 32%. Their model incorporated real-time sensor data and was validated through comprehensive statistical error metrics such as Root Mean Square Error (RMSE) and Mean Absolute Error (MAE). The study underscored the adaptability of ANNs in smart buildings by integrating HVAC systems with predictive controls to balance energy consumption and occupant comfort (Boutahri & Tilioua, 2024). Similarly, Park & Woo investigated feature selection methods combined with ANNs to optimize PMV predictions. By utilizing Principal Component Analysis (PCA) and Best Subset selection, the research pinpointed the most influential variables for accurate and efficient PMV computation. The study achieved an impressive 89.7% accuracy, demonstrating the efficacy of ANNs in reducing computational loads while maintaining predictive precision (Park & Woo, 2023).

Pavirani et al. (2024) proposed a hybrid approach by integrating PiNNs with Monte Carlo Tree Search (MCTS) algorithms. This innovative combination enabled real-time control of residential heating systems while maintaining thermal comfort and reducing energy costs. The PiNNs incorporated physical constraints into the neural network model, offering a 7% improvement in thermal comfort and a 4% reduction in energy costs compared to traditional black-box neural networks (Pavirani et al., 2024). De la Hoz-Torres et al. 2024 applied ANNs to develop adaptive thermal comfort models for naturally ventilated educational buildings. Their research demonstrated the superiority of ANN-based models over traditional PMV indices by achieving higher accuracy in thermal sensation predictions. The models were calibrated using data from a year-long monitoring campaign, revealing the significant role of adaptive behaviors in thermal comfort optimization. Lastly, Chegari et al. utilized ANNs within a surrogate-model framework to design nearly zero-energy buildings (NZEBs). This multi-objective optimization approach focused on enhancing thermal comfort and energy self-sufficiency, achieving an average improvement of 50% in comfort metrics. The surrogate model reduced computational requirements while maintaining robust performance across diverse climatic zones (Chegari et al., 2022). These studies collectively demonstrate the transformative potential of ANNs in advancing thermal comfort and energy optimization strategies. The detailed outcomes of these studies are given in Table 4.

Comparative Analysis of Energy Optimization and Comfort

The reviewed studies employing supervised machine learning techniques highlight their significant contributions to enhancing energy optimization and maintaining thermal comfort across diverse building typologies and climates. By comparing the methodologies and results across 18 selected papers, key insights can be drawn into the effectiveness and adaptability of these approaches. Tree-based models such as Random Forest (RF) and Gradient Boosted Decision Trees (GBDT) demonstrated a strong capacity for energy optimization, particularly in scenarios requiring robust feature selection and high interpretability. For example, Aparicio-Ruiz et al. showed a 5% gain in accuracy for TSV prediction in Mediterranean offices with RF by extending the variable set to 21 environmental and occupant parameters, including CO2 and running mean

Table 4. Summary of reviewed articles employing neural network applications

Study	Neural Network Type	Application	Key Metrics	Outcomes
(Boutahri & Tilioua, 2024)	ANN	PMV prediction for HVAC optimization	Accuracy (96.7% R^2), RMSE	Improved thermal comfort and energy savings in smart buildings.
(Park & Woo, 2023)	ANN with PCA	PMV dimension reduction and prediction	89.7% Accuracy, PCA Analysis	Enhanced prediction speed and accuracy by selecting key PMV parameters.
(Pavirani et al., 2024)	Physics-informed NN	Demand response and heating control	-32% MAE, -4% energy cost	Effective control with reduced computational demands using PiNN.
(de la Hoz-Torres et al., 2024)	ANN	Adaptive thermal comfort in classrooms	Improved accuracy over PMV, enhanced PMV (ePMV)	Adaptive models better suited for naturally ventilated educational buildings.
(Chegari et al., 2022)	ANN	Surrogate model for NZEB design	50% improvement in comfort metrics	Multi-objective optimization enhanced thermal comfort and energy efficiency.

temperature (Aparicio-Ruiz et al., 2023). Similarly, Brik et al. integrated RF with IoT networks to provide real-time monitoring and control, reporting 96% prediction accuracy and an 85% improvement in indoor parameter adjustment, which translated into faster restoration of comfort after disturbances (Brik et al., 2022). Lu et al. showed that RF outperformed CNN, LSTM, SVM, radial basis function (RBF) and genetic algorithm–backpropagation (GA-BP) in a hybrid radiant floor/fan-coil cooling testbed, reducing MSE by 82%, MAE by 43%, and MAPE by 68% compared with other algorithms while maintaining $R>0.99$ (Lu et al., 2024). RF also proved highly scalable in naturally ventilated schools when combined with class weighting and low-cost sensor inputs (Miao et al., 2023) highlighting its robustness under constrained data regimes. These models distinguish themselves in scenarios requiring immediate decision-making, such as hybrid cooling systems and educational buildings, by balancing energy savings with real-time thermal comfort adjustments. However, they are less suited to highly complex datasets with dynamic, non-linear interactions, as these require more adaptive learning techniques.

Regression-based models, while simpler, provided useful information for the linear relationships between environmental variables and thermal comfort indices like PMV and TSV, especially under controlled conditions. Abdellatif et al. utilized MLR with a genetic optimizer, forecasting indoor heating with lower than 1% error (adjusted $R2\approx0.9$) and achieving 43% improvement in thermal comfort over a conventional strategy (Abdellatif et al., 2022). Kumar & Kurian's (2023) Bayesian-optimized regression achieved rapid, real-time PMV and SET predictions and temperature-setpoint control, delivering measurable energy savings and higher user satisfaction in HVAC offices. Park et al. (2024) demonstrated in a field test that integrating mean radiant temperature into a thermal comfort controller yielded a 60% reduction in PMV unmet hours and more than 20% energy savings. Yet, the inherent simplicity of regression-based approaches limited its performance in contexts with large adaptive variability or strong non-linear effects, such as naturally ventilated buildings or heritage structures, where RMSEs and comfort gains lagged behind tree-based or ANN approaches.

Neural network applications stood out for their adaptability and precision in handling complex, non-linear datasets. This makes them highly effective in real-time thermal comfort control. For instance, Boutahri & Tilioua achieved a R2 of 96.7% and significant energy savings (nearly 32%) using ANNs for PMV forecasting in smart buildings, demonstrating significant energy savings without compromising occupant comfort (Boutahri & Tilioua, 2024). De la Hoz-Torres et al. created adaptive ANN comfort models for naturally ventilated educational buildings that outperformed traditional methods by integrating real-time environmental and occupant data (de la Hoz-Torres et al., 2024). Chegari et al.'s ANN surrogate model for nearly zero-energy buildings improved thermal comfort metrics by nearly 50% while reducing energy demand substantially (Chegari et al., 2022), and Pavirani et al. (2024) showed that a physics-informed neural network (PiNN) coupled with Monte Carlo Tree Search produced 32% lower MAE in thermal forecasting, 7% comfort improvement and 4% energy cost reduction over a black-box NN (Pavirani et al., 2024). Park & Woo (2023) further demonstrated that combining PCA and best-subset selection with ANN achieved 89.7% accuracy on reduced-dimension PMV, cutting computational load without sacrificing predictive precision. These examples show that, although computationally demanding, neural networks excel in contexts requiring adaptation to occupant behavior, dynamic climates and multi-objective energy–comfort balance.

Cross-model comparisons across the 18 studies reveal a clear hierarchy of suitability under different operational and climatic conditions. Tree-based ensembles (RF, GBDT, LightGBM) provide the most acceptable balance of predictive accuracy, interpretability and computational efficiency in feature-rich but moderately dynamic settings, particularly hybrid HVAC systems, IoT deployments and digital-twin maintenance frameworks. For example, RF reduced MSE by 82 % and MAE by 43% in a hybrid radiant-floor/fan-coil system (Lu et al., 2024) and maintained $R2>0.99$, while class-weighted RF models in naturally ventilated schools achieved robust performance without dense sensor networks (Miao et al., 2023). These results indicate that tree-based methods are especially advantageous where real-time decisions must be combined with variable importance screening and low-latency response. Regression-based approaches, though inherently linear, remain valuable when data complexity is low and model transparency is essential. They excel in controlled HVAC settings, offering fast convergence and interpretable parameters, as demonstrated by Abdellatif et al. (less than 1% forecast error and 43% comfort improvement) and Park et al. (2024) (60% fewer unmet PMV hours and more than 20% energy savings) in field trials. However, the comparative RMSE and comfort gains in naturally ventilated or heritage contexts consistently lag behind tree-based or ANN approaches, underscoring their limited capacity to model adaptive occupant behavior and multi-factor interactions. Neural networks and hybrid physics-informed variants clearly deliver the lowest RMSE and the highest comfort gains in real-time, adaptive contexts such as naturally ventilated schools, heritage dwellings and demand-response heating control. Boutahri & Tilioua reported R2 of nearly 0.97 and approximately 32% of energy savings; Chegari et al. achieved about 50% improvement in comfort metrics in nearly zero-energy buildings; and Pavirani et al. (2024) demonstrated 32% lower MAE, 7%

comfort gain and 4% cost reduction with PiNN-based control. Park & Woo (2023) further showed that feature selection (PCA, Best Subset) enables ANNs to reach nearly 90% accuracy with reduced input dimensions, alleviating computational burden while retaining predictive power.

To further clarify these performance relationships, a visual synthesis was developed based on representative quantitative outcomes reported across the reviewed studies. Figure 2 provides a comparative summary of these results, emphasizing the relative patterns of error reduction, energy savings, and comfort accuracy among tree-based, regression-based, and neural-network approaches. The figure presents an indicative comparison, as it reflects the characteristic performance ranges drawn from key representative works (Aparicio-Ruiz et al., 2023; Bai et al., 2022; Brik et al., 2022; Lu et al., 2024; Miao et al., 2023; Abdellatif et al., 2022; Kumar & Kurian, 2023; Park et al., 2024; Boutahri & Tilioua, 2024; Chegari et al., 2022; Pavirani et al., 2024; de la Hoz-Torres et al., 2024). Displayed values illustrate the relative magnitude of error reduction, energy savings, and comfort accuracy reported in the literature.

Collectively, these findings show that ANNs outperform tree-based and regression methods when non-linearity, occupant adaptation and multi-objective energy–comfort balance dominate. Across all model types, integrating real-time occupant behavior and environmental data emerged as the single strongest predictor of stable accuracy and energy savings. Studies omitting such inputs consistently reported higher errors or weaker generalization. This synthesis demonstrates that supervised learning approaches are not interchangeable but rather scenario-specific tools:

- RF and LightGBM excel in rapid, interpretable decisions with heterogeneous data streams;
- regression-based models are suited to stable, quasi-linear regimes;
- ANNs or PiNNs are indispensable for non-linear, occupant-centered and multi-objective optimization contexts.

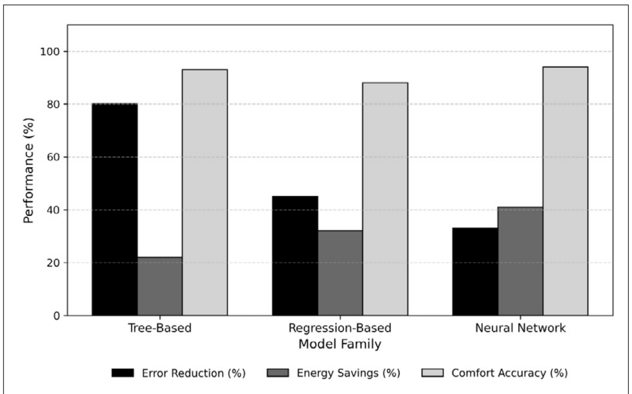

Figure 2. Comparative performance synthesis of supervised learning model families.

Identified Research Gaps

While supervised machine learning techniques have demonstrated significant potential in advancing energy optimization and thermal comfort within the built environment, several research gaps remain, limiting their widespread application and effectiveness. The review of selected studies reveals critical areas that require further investigation to address current limitations and advance the state of the art.

Integration of Real-Time Occupant Behavior

Many studies, particularly those employing regression-based and tree-based models, fail to fully integrate real-time occupant behavior into their predictive frameworks. For instance, models developed by Bai et al. and Aparicio-Ruiz et al. relied heavily on static environmental parameters, overlooking dynamic behavioral patterns such as adaptive actions or occupancy changes (Aparicio-Ruiz et al., 2023; Bai et al., 2022). This gap suggests the need for models that incorporate occupant interactions with their environments, particularly in naturally ventilated or mixed-mode buildings. To move the field forward, future work should (i) develop standardized behavior taxonomies (e.g., window/door operations, clothing adjustment, local fan/heater use) and minimal sensing protocols that can be replicated across buildings; (ii) fuse occupant-event streams with IEQ data for sequence-aware models (e.g., RF/LightGBM with lag features; LSTM/Temporal CNNs; hybrid PiNNs) and report incremental error reduction attributable to behavior; and (iii) publish ablation studies that quantify how much each behavior class improves RMSE/MAE and energy-comfort trade-offs. Such studies would directly test, in the same manner as Miao et al. (2023) and Park et al. (2024), whether adding behavior signals yields statistically significant gains over environment-only baselines.

Limited Focus on Diverse Climates and Building Typologies

The studies predominantly address specific climates or building types, such as Mediterranean climates (Aparicio-Ruiz et al., 2023) or educational buildings (Miao et al., 2023). Few have extended their applications to a broader range of climates or typologies, such as heritage buildings or high-performance green buildings. This limitation restricts the generalizability and scalability of the findings, underlining the need for research exploring diverse climatic and architectural contexts. A community benchmark of multi-climate, multi-typology datasets (e.g., classrooms, offices, heritage dwellings, NZEBs) with harmonized labels (PMV/TSV/ePMV) and common splits for external validation is recommended. Protocols should require reporting per-climate and per-typology performance, enabling fair cross-study comparisons similar to Liu & Ma, 2023 and Xi et al. (2024). Model cards should include “applicability statements” that explicitly state validated Köppen–Geiger zones and building archetypes.

Data Scarcity and Model Robustness

Several studies highlighted the challenges of acquiring high-quality, comprehensive datasets. For instance, Hosamo et al. 2023 and Brik et al. 2022 relied on IoT networks, which, while effective, are resource-intensive and not universally accessible. Additionally, many models were tested on limited datasets, raising concerns about their robustness and applicability in real-world scenarios. Future studies should focus on developing models that are robust to incomplete or noisy data and leverage innovative data augmentation techniques. Concrete next steps include: (i) Adopting nested cross-validation and leakage-safe feature selection to improve reliability across all three families (as issues were noted in multiple papers); (ii) stress-test models under missingness and sensor drift; and (iii) employing transfer learning/domain adaptation between climates and building types (e.g., training in Mediterranean offices and adapting to NV schools) with explicit reporting of adaptation gains. Open baselines should include LightGBM/RF, linear models, and at least one ANN to anchor robustness claims.

Evaluation practice deficiencies

Across the corpus, nested cross-validation and external validation were uncommon. When hyper-parameter tuning, feature selection, or preprocessing (imputation, scaling, resampling) are performed outside a nested scheme or on the full dataset, information leakage can inflate accuracy and understate uncertainty. Likewise, evaluating only on the same sites or periods used for development risks optimism and weakens claims of generalizability. Minimum leakage-safe practice should include repeated nested k-fold CV (all modeling operations confined to inner folds) and an external test via site-out or time-split protocols (different buildings, seasons, or terms). Reporting should add calibration metrics, fold-wise variance, and failure modes observed under robustness checks. Adopting these standards will immediately improve the credibility and comparability of results across model families.

Real-Time Processing and Adaptability

Neural network-based studies, such as those by Boutahri & Tilioua and De la Hoz-Torres et al., demonstrated strong adaptability but often required substantial computational resources (Boutahri & Tilioua, 2024; de la Hoz-Torres et al., 2024). These models struggle with real-time processing in low-resource environments, particularly in remote or economically constrained areas. Addressing this limitation by optimizing algorithms for computational efficiency or leveraging edge computing could make these methods more accessible and practical. Future research should (i) benchmark latency, memory, and power on representative edge hardware; (ii) evaluate model compression (quantization/pruning/knowledge distillation) and feature reduction (as in Park & Woo, 2023) with comfort/energy accuracy retained; and (iii) report end-to-end control-loop stability (response time to setpoint changes, overshoot/undershoot) alongside prediction metrics.

Beyond raw accuracy, the real-world applicability of ANN/PiNN approaches is shaped by total cost of ownership and operational risk. Training often depends on specialized accelerators and curated pipelines, while inference on site can exceed the latency, memory, and power envelopes of legacy BMS or low-cost edge controllers; cloud off-loading adds recurring costs, privacy/compliance concerns, and network fragility. Scaling across buildings also requires site-specific calibration and continuous monitoring for drift, with nontrivial data quality checks, re-training, versioning, and rollbacks. Limited transparency can slow operator troubleshooting when comfort or IAQ alarms trigger, reducing trust compared with simpler, interpretable controllers. A pragmatic stance is to prefer compact tree ensembles or linear controllers when ANN gains are marginal or budgets are constrained, reserving ANN/PiNN solutions for strongly non-linear, occupant-adaptive contexts where demonstrated energy/comfort benefits outweigh compute and maintenance costs. When ANNs are deployed, studies should include simple guardrails on set-point changes, document fail-safe modes for sensor/connectivity faults, and report measured latency, memory, and power for the compressed/distilled model on the target edge device to demonstrate field readiness.

Balancing Energy Savings with Thermal Comfort

While energy optimization is an important focus, few studies explicitly quantify the trade-offs between energy savings and thermal comfort. For example, the works by Kumar & Kurian (2023) and Mousavi et al. (2023) emphasized energy savings but provided limited insights into how these savings impact occupant comfort under varying conditions. Future research should aim to establish a clearer balance between these objectives, incorporating adaptive comfort models that prioritize human well-being without significant energy penalties. For explicit multi-objective formulations with Pareto fronts (comfort vs. kWh/cost), reporting dominated vs. non-dominated solutions and sensitivity to seasonal/occupancy regimes is beneficial. Studies like Pavirani et al. (2024) and Chegari et al. (2024) provide templates; forthcoming work should standardize comfort violation metrics (e.g., unmet PMV/TSV hours, ePMV bands) and quantify comfort "cost" per unit energy saved.

Integration with Emerging Technologies

Most studies reviewed did not explore the integration of supervised learning techniques with emerging technologies such as digital twins, advanced IoT frameworks, or blockchain for data security and decentralization. The work by Hosamo et al. (2023) on digital twins stands as a notable exception but highlights the potential for combining machine learning with advanced technologies to enhance predictive accuracy, energy efficiency, and comfort management. Future work should couple calibrated digital twins with supervised learning for fault diagnostics and proactive control, evaluating whether twin-in-the-loop supervision reduces failure rates and comfort violations beyond RF-only baselines (as hinted by Hosamo et al., 2023). Data governance should be addressed via privacy-preserving pipelines (federated learning, differential privacy) to enable cross-site generalization without sharing raw occupant data.

Long-Term Validation Studies

Many studies evaluated their models using short-term datasets or simulations, with limited validation in real-world, long-term operational settings. For example, Lu et al. (2024) demonstrated energy savings in hybrid cooling systems but lacked long-term empirical data to substantiate these findings under varying operational conditions. Longitudinal studies that track model performance over extended periods are needed to validate their reliability and effectiveness. More than 12-month deployments spanning seasons and occupancy cycles, with pre-registered analysis plans, drift detection, and periodic re-calibration rules are recommended. Reports should include durability of gains (R2/MAE stability, comfort violations, energy bills) and failure mode analyses (sensor outages, occupancy anomalies).

Equity and Accessibility Considerations

A recurring gap is the lack of focus on making these technologies accessible in economically constrained or developing regions. Models relying on high-cost infrastructure, such as IoT networks or advanced computational systems, are less applicable in these settings. Research aimed at creating cost-effective and scalable solutions, like the RF-based model by Miao et al. (2023), could address this inequity. Priorities include low-cost sensing kits, sparse-feature models that maintain accuracy with minimal inputs, and edge-deployable controllers. Studies should report a "cost-to-accuracy" curve and provide open designs/bills of materials so that public schools and small offices can reproduce the results.

In sum, the empirical patterns across the 18 studies suggest a pathway for targeted progress: (1) Add behavior signals and temporal structure to supervised models; (2) validate across climates/typologies using shared benchmarks; (3) enforce leakage-safe evaluation (nested CV, external tests) and robustness checks; (4) operationalize real-time constraints on edge hardware; (5) optimize explicitly on the comfort-energy Pareto frontier; (6) integrate digital-twin supervision and privacy-preserving data pipelines; (7) extend evaluations to multi-season deployments; and (8) prioritize low-cost, reproducible solutions. Addressing these items will convert today's promising but fragmented results into generalizable, field-ready ML frameworks that reliably balance energy efficiency and occupant comfort.

Strategies for Progress: Practical Implications

The limitations identified in Section 3.3 are not isolated shortcomings but stem from recurring structural and methodological challenges within the current research landscape. Recognizing the underlying reasons for these shortcomings and proposing practical strategies to address them can accelerate progress in supervised machine learning for thermal comfort and energy efficiency. A first and persistent limitation arises from the restricted data coverage and quality of existing studies. Many models are developed from single buildings, limited climates, or short monitoring periods. This narrow scope limits the diversity of environmental conditions, occupant behaviors, and building typologies captured in the datasets. As a result, models often lack external validity and show performance drops when applied to new settings. Addressing this gap requires coordinated efforts to build multi-site and multi-season datasets with harmonized comfort indices, consistent sensor metadata, and clear contextual variables such as occupancy schedules or ventilation strategies. When such large-scale data collection is not feasible, researchers can still improve reliability through nested cross-validation, leakage-safe feature selection, and data augmentation or simulation of unobserved conditions. By stress-testing models under missing data or sensor drift, studies can quantify robustness before deployment.

Another major cause of current shortcomings is the limited integration of occupant behavior into predictive frameworks. Most models rely heavily on environmental variables and treat occupants as passive recipients of indoor conditions. Yet evidence from adaptive comfort research shows that actions such as window opening, clothing adjustment, or use of local fans can substantially shift comfort thresholds. A key strategy is to develop standardized, low-cost protocols for capturing occupant actions, either through simple binary sensors or self-reports linked to time-stamped environmental data. These behavioral event streams can then be incorporated into supervised models as lagged or sequential features, or through temporal and sequence-aware architectures such as LSTMs, temporal CNNs, or hybrid physics-informed neural networks. Publishing ablation studies that explicitly compare environment-only models with behavior-enhanced models would help quantify the added value and set benchmarks for future work.

A third limitation stems from generalizing across climates and building typologies. Most models have been validated only in Mediterranean offices, naturally ventilated schools, or similar narrow archetypes. This raises the risk that models encode climate-specific correlations rather than universal principles. Researchers can overcome this by developing and sharing multi-climate benchmark datasets with fixed training–testing splits, and by reporting performance separately for each climate zone and building archetype.

Domain adaptation and transfer-learning methods can be tested explicitly (for example, training on Mediterranean offices and adapting to educational buildings in temperate zones) with reported adaptation gains. Model "applicability statements" could then state validated Köppen–Geiger zones and archetypes, improving transparency for practitioners.

Computational constraints also play a key role. High-capacity neural networks offer excellent accuracy but may be too resource-intensive for real-time, edge-level control. Without careful attention to latency, memory, and power, these models cannot be integrated into HVAC controllers or low-cost sensing platforms. Strategies include model compression techniques such as quantization, pruning, or knowledge distillation, combined with feature-reduction approaches to lower input dimensionality while retaining predictive power. Benchmarking models on representative embedded hardware, and reporting end-to-end control loop metrics such as response time and overshoot, would make research outcomes far more actionable for practitioners.

Another widespread shortcoming is the one-sided focus on either comfort or energy without explicitly quantifying trade-offs. This obscures the true cost of achieving comfort gains or energy savings. Future work should adopt multi-objective optimization frameworks to map Pareto fronts of comfort versus energy, and employ standardized comfort violation metrics such as unmet PMV or TSV hours. Reporting comfort "cost per unit energy saved" and seasonal sensitivity analyses would help designers and operators choose balanced strategies and compare across studies. This approach transforms models from black-box predictors into decision-support tools with clear operational implications.

Finally, issues of data governance, privacy, and reproducibility constrain the cross-site validation needed for robust models. Sharing raw occupant or environmental data across institutions is often impractical or unethical. Emerging privacy-preserving methods such as federated learning or differential privacy can allow multiple sites to train a shared model without exchanging raw data. Accompanying open-source code, baseline models, and clear reporting checklists (including dataset splits, leakage tests, and calibration metrics) will further strengthen reproducibility and accelerate uptake. Together, these strategies form a coherent roadmap for converting today's promising but fragmented studies into reliable, scalable tools for building practice. By combining broader, higher-quality datasets with behavior-aware features, multi-climate benchmarking, edge-ready model designs, explicit multi-objective optimization, and privacy-preserving collaboration, the field can move beyond narrow proofs of concept to deliver field-ready, occupant-centered, energy-efficient control systems. In practical terms, this means HVAC systems capable of dynamically adapting to both environmental changes and human actions, design recommendations grounded in diverse climates and typologies, and machine learning models that can be deployed even in low-resource settings. Ultimately, such advances will help translate the theoretical potential of supervised machine learning into widespread real-world impact, supporting carbon reduction, improved occupant well-being, and the broader goals of sustainable architecture and urban development.

CONCLUSION

This study employed a semi-systematic review to examine the application of supervised learning approaches in thermal comfort prediction and energy optimization within the built environment, using a transparent and replicable search, and screening process. By categorizing the reviewed studies into tree-based models, regression-based models, and neural network applications, their unique strengths, methodological contributions, and practical applications were highlighted.

Tree-based models, such as Random Forest and Gradient Boosted Decision Trees, stand out in interpretability and feature selection, which makes them effective tools for real-time decision-making in hybrid systems and IoT-enhanced frameworks. Regression-based models, characterized by their simplicity and linear focus, are highly suited for controlled environments and scenarios requiring efficient and scalable solutions. Neural networks demonstrated good adaptability and precision, particularly in dynamic, non-linear scenarios requiring real-time adjustments, such as naturally ventilated or smart buildings. Supervised learning approaches collectively showed a substantial potential in improving building energy efficiency and occupant comfort. Neural network models, in particular, consistently delivered high accuracy and adaptability, enabling significant energy savings while maintaining or enhancing thermal comfort. However, tree-based and regression models remain valuable alternatives in contexts with constrained computational resources or data availability, providing practical and scalable solutions. Synthesizing across these strands, several practice-oriented takeaways emerge:

- In practice, selecting the model family to fit the operational context yields the best results, with tree ensembles balancing accuracy and interpretability in feature-rich yet moderately dynamic settings, regression suiting simple and transparent control, and ANNs/PiNNs excelling in strongly non-linear, occupant-adaptive scenarios.
- With appropriate configuration and validation, supervised models can deliver measurable energy savings without degrading thermal comfort.
- Incorporating real-time occupant actions and contextual variables improves generalization across climates and building typologies.

- Leakage-safe pipelines with repeated nested cross-validation plus external tests are essential for reliable, comparable claims across model families.
- Deployment constraints (compute, latency, power, maintainability) often favor compact ensembles or compressed neural networks on edge hardware, supplemented by guardrails and fail-safes for closed-loop control.

The study identified several critical research gaps and methodological limitations that constrain the broader application of these techniques. Key gaps include the limited integration of real-time occupant behavior, insufficient focus on diverse climatic and building typologies, challenges in data availability and model robustness, and the trade-offs between energy savings and occupant comfort. The lack of long-term validation studies and limited integration with emerging technologies, such as digital twins and edge computing, further underscore the need for advanced research efforts. In addition, as a semi-systematic review with a defined search scope and standardized inclusion criteria, the analysis is inherently bounded by its database coverage and screening framework. Although every effort was made to ensure transparency and replicability, the reliance on a single indexing source and the absence of a formal risk-of-bias appraisal may have led to the omission of a small number of relevant studies. This focused approach was adopted to maintain methodological consistency, clarity, and reproducibility while minimizing redundancy among overlapping indexing platforms. Future reviews may broaden the search scope as the field expands and diversifies.

Acknowledging both research gaps and methodological constraints provides a transparent basis for interpreting the findings and highlights opportunities for future work to develop more inclusive, robust, and scalable machine learning-based solutions. As a minimum good-practice standard, future studies should report repeated nested cross-validation (with all modeling operations inside folds) together with an external validation on independent sites or time windows, alongside calibration and robustness analyses. The practical implications of this research offer that machine learning-driven HVAC systems represent a transformative approach to sustainable building practices, enabling buildings to dynamically adapt to changing conditions while balancing energy efficiency and human comfort. These systems have the potential to reduce energy consumption, lower operational costs, and enhance occupant well-being, contributing to global efforts toward carbon neutrality and sustainable development. As technology and computational capabilities advance, the integration of supervised learning in HVAC systems and thermal comfort management will likely play an essential role in shaping the future of energy-efficient and human-centered architecture.

Appendix: https://jag.journalagent.com/megaron/abs_files/MEGARON-02256/MEGARON-02256_(2)_Appendix_Table_A1.pdf

ETHICS: There are no ethical issues with the publication of this manuscript.

PEER-REVIEW: Externally peer-reviewed.

CONFLICT OF INTEREST: The authors declared no potential conflicts of interest with respect to the research, authorship, and/or publication of this article.

FINANCIAL DISCLOSURE: The authors declared that this study has received no financial support.

REFERENCES

Abdellatif, M., Chamoin, J., Nianga, J. M., & Defer, D. (2022). A thermal control methodology based on a machine learning forecasting model for indoor heating. *Energy and Buildings, 255*, 111692. https://doi.org/10.1016/j.enbuild.2021.111692

Aparicio-Ruiz, P., Barbadilla-Martín, E., Guadix, J., & Nevado, J. (2023). Analysis of variables affecting indoor thermal comfort in Mediterranean climates using machine learning. *Buildings, 13*(9), 2215. https://doi.org/10.3390/buildings13092215

Arakawa Martins, L., Soebarto, V., & Williamson, T. (2022). A systematic review of personal thermal comfort models. *Building and Environment, 207,* 108502. https://doi.org/10.1016/j.buildenv.2021.108502

Bai, Y., Liu, K., & Wang, Y. (2022). Comparative analysis of thermal preference prediction performance in different conditions using ensemble learning models based on ASHRAE Comfort Database II. *Building and Environment, 223*, 109462. https://doi.org/10.1016/j.buildenv.2022.109462

Boutahri, Y., & Tilioua, A. (2024). Machine learning-based predictive model for thermal comfort and energy optimization in smart buildings. *Results in Engineering, 22*, 102148. https://doi.org/10.1016/j.rineng.2024.102148

Brik, B., Esseghir, M., Merghem-Boulahia, L., & Hentati, A. (2022). Providing convenient indoor thermal comfort in real-time based on energy-efficiency IoT network. *Energies, 15*(3), 808. https://doi.org/10.3390/en15030808

Chegari, B., Tabaa, M., Simeu, E., Moutaouakkil, F., & Medromi, H. (2022). An optimal surrogate-model-based approach to support comfortable and nearly zero energy buildings design. *Energy, 248*, 123584. https://doi.org/10.1016/j.energy.2022.123584

de la Hoz-Torres, M. L., Aguilar, A. J., Ruiz, D. P., & Martínez-Aires, M. D. (2024). An investigation of indoor thermal environments and thermal comfort in naturally ventilated educational buildings. *Journal of Building Engineering, 84*, 108677. https://doi.org/10.1016/j.jobe.2024.108677

Essamlali, I., Nhaila, H., & El Khaili, M. (2024). Supervised machine learning approaches for predicting key pollutants and for the sustainable enhancement of urban air quality: A systematic review. *Sustainability, 16*(3), 976. https://doi.org/10.3390/su16030976

Feng, Y., Liu, S., Wang, J., Yang, J., Jao, Y. L., & Wang, N. (2022). Data-driven personal thermal comfort prediction: A literature review. *Renewable and Sustainable Energy Reviews, 161*, 112357. https://doi.org/10.1016/j.rser.2022.112357

Gupta, V., & Deb, C. (2022). Energy retrofit analysis for an educational building in Mumbai. *Sustainable Futures, 4*, 100096. https://doi.org/10.1016/j.sftr.2022.100096

Halhoul Merabet, G., Essaaidi, M., Ben Haddou, M., Qolomany, B., Qadir, J., Anan, M., Al-Fuqaha, A., Abid, M. R., & Benhaddou, D. (2021). Intelligent building control systems for thermal comfort and energy-efficiency: A systematic review of artificial intelligence-assisted techniques. *Renewable and Sustainable Energy Reviews, 144*, 110969. https://doi.org/10.1016/j.rser.2021.110969

Han, X., Hu, Z., Li, C., Wu, J., Li, C., & Sun, B. (2023). Prediction of human thermal comfort preference based on supervised learning. *Journal of Thermal Biology, 112*, 103484. https://doi.org/10.1016/j.jtherbio.2023.103484

Hosamo, H. H., Nielsen, H. K., Kraniotis, D., Svennevig, P. R., & Svidt, K. (2023). Improving building occupant comfort through a digital twin approach: A Bayesian network model and predictive maintenance method. *Energy and Buildings, 288*, 112992. https://doi.org/10.1016/j.enbuild.2023.112992

Kumar, T. M. S., & Kurian, C. P. (2023). Real-time data based thermal comfort prediction leading to temperature setpoint control. *Journal of Ambient Intelligence and Humanized Computing, 14*(9), 12049–60. https://doi.org/10.1007/s12652-022-03754-8

Lala, B., & Hagishima, A. (2022). A review of thermal comfort in primary schools and future challenges in machine learning-based prediction for children. *Buildings, 12*(11), 2007. https://doi.org/10.3390/buildings12112007

Liu, H., & Ma, E. (2023). An explainable evaluation model for building thermal comfort in China. *Buildings, 13*(12), 3107. https://doi.org/10.3390/buildings13123107

Lu, S., Cui, M., Gao, B., Liu, J., Ni, J., Liu, J., & Zhou, S. (2024). A comparative analysis of machine learning algorithms in predicting the performance of a combined radiant floor and fan coil cooling system. *Buildings, 14*(6), 1659. https://doi.org/10.3390/buildings14061659

Miao, S., Gangolells, M., & Tejedor, B. (2023). Data-driven model for predicting indoor air quality and thermal comfort levels in naturally ventilated educational buildings using easily accessible data for schools. *Journal of Building Engineering, 80*, 108001. https://doi.org/10.1016/j.jobe.2023.108001

Moshood, T. D., Nawanir, G., Lee, C. K., & Fauzi, M. A. (2024). Toward sustainability and resilience with Industry 4.0 and Industry 5.0. *Sustainable Futures, 8*, 100349. https://doi.org/10.1016/j.sftr.2024.100349

Mousavi, S., Gheibi, M., Wacławek, S., Smith, N. R., Hajiaghaei-Keshteli, M., & Behzadian, K. (2023). Low-energy residential building optimisation for energy and comfort enhancement in semi-arid climate conditions. *Energy Conversion and Management, 291*, 117264. https://doi.org/10.1016/j.enconman.2023.117264

Ouzzani, M., Hammady, H., Fedorowicz, Z., & Elmagarmid, A. (2016). Rayyan—a web and mobile app for systematic reviews. *Systematic Reviews, 5*(1), 210. https://doi.org/10.1186/s13643-016-0384-4

Park, J., Kim, T., Kim, D., Alghimlas, F., AlRagom, F., Choi, H., & Cho, H. (2024). Field test of machine-learning based mean radiant temperature estimation methods for thermal comfort-integrated air-conditioning control improvement and energy savings. *Energy Reports, 11*, 5682–702. https://doi.org/10.1016/j.egyr.2024.05.040

Park, K. Y., & Woo, D. O. (2023). PMV dimension reduction utilizing feature selection method: Comparison study on machine learning models. *Energies, 16*(5), 2419. https://doi.org/10.3390/en16052419

Pavirani, F., Gokhale, G., Claessens, B., & Develder, C. (2024). Demand response for residential building heating: Effective Monte Carlo Tree Search control based on physics-informed neural networks. *Energy and Buildings, 311*, 114161. https://doi.org/10.1016/j.enbuild.2024.114161

Qavidel Fard, Z., Zomorodian, Z. S., & Korsavi, S. S. (2022). Application of machine learning in thermal comfort studies: A review of methods, performance and challenges. *Energy and Buildings, 256*, 111771. https://doi.org/10.1016/j.enbuild.2021.111771

Sibyan, H., Svajlenka, J., Hermawan, H., Faqih, N., & Arrizqi, A. N. (2022). Thermal comfort prediction accuracy with machine learning between regression analysis and naïve Bayes classifier. *Sustainability, 14*(23), 15663. https://doi.org/10.3390/su142315663

Wolff, R. F., Moons, K. G. M., Riley, R. D., Whiting, P. F., Westwood, M., Collins, G. S., Reitsma, J. B., Kleijnen, J., & Mallett, S. (2019). PROBAST: A tool to assess the risk of bias and applicability of prediction model studies. *Annals of Internal Medicine, 170*(1), 51–8. https://doi.org/10.7326/M18-1376

Xi, H., Wang, B., & Hou, W. (2024). Machine learning-based prediction of indoor thermal comfort in traditional Chinese dwellings: A case study of Hankou Lifen. *Case Studies in Thermal Engineering, 61*, 105048. https://doi.org/10.1016/j.csite.2024.105048

Zhang, W., Wu, Y., & Calautit, J. K. (2022). A review on occupancy prediction through machine learning for enhancing energy efficiency, air quality and thermal comfort in the built environment. *Renewable and Sustainable Energy Reviews, 167*, 112704. https://doi.org/10.1016/j.rser.2022.112704

Zhou, S. L., Shah, A. A., Leung, P. K., Zhu, X., & Liao, Q. (2023). A comprehensive review of the applications of machine learning for HVAC. *DeCarbon, 2,* 100023. https://doi.org/10.1016/j.decarb.2023.100023