Indoor Wellness Optimization in Smart Buildings Based on Ambient Variable Dimensionality Reduction

Gaetano Ingenito, Oumaima Afif, *Graduate Student Member, IEEE*, Aldo Romani, *Member, IEEE*, and Marco Tartagni, *Member, IEEE*

***Abstract*—Achieving thermal comfort in indoor environments is crucial for both occupant well-being and energy sustainability. Although recent machine learning (ML) and deep reinforcement learning (DRL) approaches have shown promise for real-time heating, ventilation, and air conditioning (HVAC) control, their reliance on simulated environments, high computational costs, and limited generalizability hinder practical adoption. In particular, these methods often require large datasets and offer limited interpretability of variable correlations. This article proposes a lightweight and scalable alternative by integrating multivariate analysis (MVA) with the Internet of Things (IoT) data streams for real-time comfort prediction and control on a flexible number of variables. The approach combines a nonlinear locally weighted regression (LWR) model with a sequential quadratic programming (SQP) optimization layer, enforcing physical and operational constraints to produce interpretable and feasible control actions. The optimization strategy is highly flexible, allowing independent boundary conditions on environmental variables. Dimensionality reduction techniques enable deployment on resource-constrained hardware without sacrificing accuracy. The proposed framework is validated on a real-world dataset collected from three indoor zones at the University of Bologna (Campus of Cesena), comprising 765 observations of environmental variables, including temperature (Temp), relative humidity (RH), and indoor air quality (IAQ), sampled every 20 min, along with hourly outdoor conditions and subjective satisfaction data from occupant surveys. The results show predicted wellness improvements of 14.4%, 77.8%, and 89.6% in three different rooms.**

***Index Terms*—Digital twins, environmental monitoring, indoor wellness, Internet of Things (IoT), machine learning (ML), multivariate analysis (MVA).**

I. Introduction

ACHIEVING optimal thermal comfort is becoming increasingly important in the pursuit of sustainable and human-centered built environments [1], [2]. Thermal comfort, shaped by environmental, physiological, and psychological factors, directly impacts productivity, health, and overall well-being within indoor spaces [3], [4], [5].

Recent advancements in real-time indoor environmental monitoring and control have been driven by the rise of machine learning (ML), particularly for temperature (Temp) regulation and occupant satisfaction. ML-augmented heating, ventilation, and air conditioning (HVAC) systems, leveraging time-series data, have enabled dynamic prediction and control strategies. Notably, the supervised learning and deep reinforcement learning (DRL) have gained traction for their ability to personalize comfort and improve energy efficiency [6], [7].

However, despite their promise, DRL and other ML-based approaches often suffer from key limitations: they are computationally intensive, frequently rely on simulation data, and lack generalizability to diverse real-world conditions. For example, Al-Adhami et al. [8] and Alsaleem et al. [9] proposed the real-time and personalized thermal comfort models, yet both approaches impose significant system complexity and demand substantial computational resources. Similarly, the advanced algorithms like deep deterministic policy gradient (DDPG) have shown improved accuracy and energy efficiency [7] but remain difficult to deploy at scale due to their reliance on simulated environments and computational overhead. Moreover, the hybrid methods, such as those combining DRL and model predictive control (MPC) [10], demonstrate improved performance in simulation, but their complexity and dependency on tools like EnergyPlus [11] hinder real-time adoption in actual HVAC systems.

Most ML-based approaches for HVAC control deliver promising results in simulation environments but are rarely adopted in real-world building operations [12]. This gap between model development and field implementation is due to three persistent challenges. First, installation complexity: integrating an ML or reinforcement learning controller with legacy building management system (BMS) infrastructure, often based on protocols like BACnet or Modbus, requires custom software bridges, safety interlocks, and several days of on-site configuration [5], [13]. Second, deployment cost: additional sensors, edge computing devices, and secure data transmission pipelines may offset potential energy savings, especially in retrofit scenarios [5], [14]. Third, limited interpretability: building operators are hesitant to adopt closed-box

Received 21 March 2026; revised 23 May 2026 and 15 June 2026; accepted 29 June 2026. Date of publication 3 July 2026; date of current version 26 August 2026. This work was supported in part by the Agritech National Research Center through European Union Next-GenerationEU (PNRR)–Mission 4 Component 2, Investment 1.4–D.D. 1032 17/06/2022, under Grant CN00000022; and in part by the Chips Joint Undertaking (Chips JU) and its Members under Grant 101112286 (LoLiPoP IoT). *(Gaetano Ingenito and Oumaima Afif are co-first authors.) (Corresponding author: Oumaima Afif.)*

The authors are with the Advanced Research Center on Electronic Systems (ARCES), University of Bologna, 47522 Cesena, Italy (e-mail: gaetano.ingenito@studio.unibo.it; oumaima.afif2@unibo.it; aldo.romani@unibo.it; marco.tartagni@unibo.it).

Digital Object Identifier 10.1109/JIOT.2026.3710119

© 2026 The Authors. This work is licensed under a Creative Commons Attribution 4.0 License.
For more information, see https://creativecommons.org/licenses/by/4.0/

control policies whose internal logic cannot be easily understood or verified.

For instance, study [15] developed a hybrid architecture combining a convolutional neural network-based recurrent unit Temp forecasting model with a stacked ensemble (XGBoost, random forest, and deep neural network) for predicted mean vote (PMV) estimation. While this architecture improves predictive accuracy compared to single-model baselines, it introduces considerable computational demand and depends on the availability of large, labeled datasets conditions, which are uncommon in most residential and small commercial buildings.

Finally, Xue et al. [16] proposed an inverse-PMV set-point estimator that eliminates iterative optimization by directly learning the mapping from a target PMV value to the corresponding room Temp. The PMV, originally developed by Fanger [17], is an index that predicts the average thermal sensation of a large group of people on a seven-point scale, based on the human body's heat balance. This simplification significantly reduces computational overhead and supports real-time applicability. However, their validation relies solely on TRNSYS/MATLAB co-simulation, without incorporating the Internet of Things (IoT)-based sensing, occupant variability, or hardware-in-the-loop testing. Given these challenges, it becomes crucial to explore lighter, more adaptable alternatives that maintain predictive performance while supporting real-time control in embedded systems.

The approach explored in this work is to employ variable dimensionality reduction techniques within the multivariate analysis (MVA) framework, which have been successfully applied in sensing technologies, as an alternative, or better, as a complementary strategy, to deep-learning methods [18], [19], [20]. Specifically, for the task addressed in this article, we propose combining classical constrained optimization techniques, such as sequential quadratic programming (SQP), with a reduced-variable representation obtained from MVA. In addition, the compact multivariate time-series models, such as state-space locally weighted regression (LWR) or partial least square (PLS) regressions, can capture the joint effects of weather, occupancy, and equipment states while remaining data-efficient and inherently interpretable. Embedding these models within an SQP layer further enforces hard constraints and actuator limitations, thus providing operators with clear bounds and transparent decision logic.

Within this context, the principal component analysis (PCA)-based latent-variable methods and local models have already been explored in building applications, but mostly as separate tools. PCA-type methods are commonly used for HVAC monitoring, fault detection, and operating-state assessment [21], [22], while local models have been used for personalized comfort prediction or HVAC dynamics modeling [23], [24]. Recent building-control studies also emphasize that interpretability and deployability are critical for practical adoption [12], [25]. However, these elements are rarely combined into a single actionable pipeline. This work addresses this gap by using the dynamic score space not only for dimensionality reduction but also as the domain where local satisfaction prediction and constrained feasible optimization are performed.

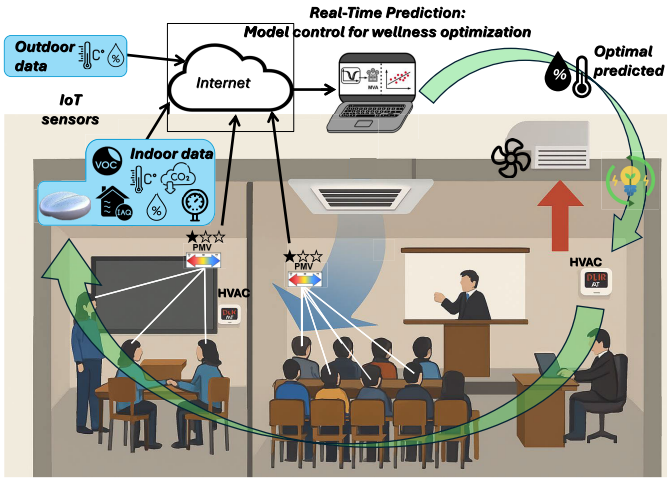

Fig. 1. Big picture of the proposed approach.

To this end, this article proposes an MVA-based framework integrated with IoT technologies to enable real-time prediction and control of indoor wellness factors, as shown in Fig. 1. The connected sensors, network services and infrastructures, and user feedback via personal devices are key enablers for this approach, supporting different phases of operation from model training to prediction. From a data analysis perspective, unlike traditional ML and deep learning (DL) approaches, the proposed method incorporates time dependencies through dynamic data modeling while avoiding resource-heavy computations. Specifically, dynamic PCA (dPCA) is employed to capture both the autocorrelation and cross correlation among environmental variables over time, aspects that conventional PCA often fails to address due to its static nature.

One of the main motivations for employing dPCA in this study is its ability to mitigate nonlinearity in the data, something traditional static PCA struggles to handle effectively. While standard PCA excels at removing linear correlations among variables, it often fails to account for nonlinear relationships and time dependencies that are common in environmental monitoring contexts. dPCA overcomes this limitation by incorporating lagged variables, allowing the model to better capture temporal patterns and hidden structures that static methods may overlook. This approach effectively reduces spurious correlations caused by time-lagged effects and improves interpretability in the presence of nonlinear behaviors. This choice is further supported by the findings of our previous work [26], where nonlinear models consistently outperformed linear ones in predicting user satisfaction from environmental data. Together, these considerations strongly justify the adoption of dPCA in the proposed framework. A nonlinear LWR model is therefore used as a baseline for multidimensional environmental optimization. Designed to operate on resource-constrained microcontrollers, this approach supports scalable and energy-efficient deployment.

The experimental validation is conducted on a real-world dataset collected from three rooms, as shown in Fig. 2, at the

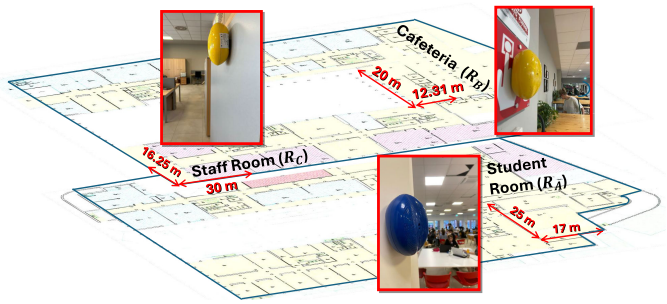

Fig. 2. On-site pictures of the INAIR devices on a simplified map of the Cesena campus, R_A on the ground floor; R_B and R_C on the first floor.

University of Bologna (Campus of Cesena), each characterized by different occupancy levels—namely, high, intermediate, and low. The dataset was previously introduced in [26] but is reprocessed here using new preprocessing techniques, by removing irrelevant variables, mean-centering, and autoscaling to standardize the data. In particular, the reprocessing introduced in this work includes a systematic data-quality screening stage to detect and remove rare spurious samples (mainly affecting gas-based IAQ channels) caused by sensor handling or transient contamination.

The monitoring campaign was conducted from November 23 at 09:00 to January 13 at 18:00. Only working-hour samples, from 09:00 to 18:00, were retained at 20-min resolution, and the holiday on December 8 was excluded. After synchronization, filtering, and outlier removal, the final dataset consists of 765 synchronized observations. Since the campaign covered only a late-autumn/winter period, multiseason generalization is not claimed. The monitoring campaign was designed to capture heterogeneous occupancy and usage patterns across three indoor environments while minimizing intrusiveness. Two sensing nodes were deployed: sensor S_1 was relocated weekly between Room R_A and Room R_B, whereas sensor S_2 remained fixed in Room R_C throughout the acquisition period. The monitored spaces differ substantially in geometry and occupancy dynamics. Room R_A (student room, 425 m^2) and Room R_B (cafeteria, 246.2 m^2) exhibit highly variable usage with peak occupancies exceeding 100 individuals, while Room R_C (staff room, $\approx$487.5 m^2) is characterized by more stable and sparse occupancy (typically $\approx$26 people). This diversity provides distinct operational regimes for evaluating the proposed framework under different usage profiles.

The subjective comfort-related labels were collected through structured questionnaires distributed via messaging applications. In total, 46 participants contributed responses in R_A, 27 in R_B, and 16 in R_C. The analysis uses aggregated responses to reduce individual-level variability.

Sensor data are collected from the three distinct rooms and augmented with external Temp and relative humidity (RH) measurements retrieved via the Visual Crossing API [27]. The external weather variables were resampled to the same 20-min timestamps as the indoor sensors; a timestamp-consistency check confirmed exact indoor–outdoor alignment after preprocessing (zero time offset and identical sample counts per room). In addition to objective environmental parameters, subjective comfort metrics, including the PMV and user-reported satisfaction scores, are included in the dataset.

TABLE I
MODEL- AND OPTIMIZATION-RELATED NOTATION

Symbol	Definition
$\boldsymbol{X}$	The augmented dataset with lagged observations
$\boldsymbol{X}_{\text{ref}}$	The reference subset with high satisfaction
$\boldsymbol{X}_{\text{obj}}$	The objective subset to be optimized
$\boldsymbol{T}$	The scores of dataset matrix $\boldsymbol{X}$
$\boldsymbol{T}_{\text{ref}}$	The reference score point
$\boldsymbol{T}_{\text{obj}}$	The objective score point before optimization
$\boldsymbol{T}'_{\text{obj}}$	The objective score point after applying $\boldsymbol{\Delta}$
$\boldsymbol{T}'$	The single score point
$\boldsymbol{L}$	The loading of dataset matrix $\boldsymbol{X}$
$\boldsymbol{K}$	The number of variables
$\boldsymbol{N}$	The number of observations
$\boldsymbol{A}$	The number of principal components (PCs)
$\boldsymbol{M}$	The total number of recorded variables
$\boldsymbol{P}$	The transformed dataset of PCA on the augmented $\boldsymbol{X}$
S	The singular value decomposition (SVD) of $\boldsymbol{P}$
r	The # of S
$\boldsymbol{\Delta}$	The variation vector
$\boldsymbol{\delta}_k$	The search direction at iteration k
$\mathcal{L}$	The Lagrangian function
∇, ∇^2	The gradient and Hessian of $\mathcal{L}$
$f(\boldsymbol{x})$	The objective function to minimize
$g(\boldsymbol{x})$	Equality constraint function
$h(\boldsymbol{x})$, $k(\boldsymbol{x})$	Inequality constraint functions
s, t	Slack variables for inequality handling
λ, θ, ϕ	Lagrange multipliers
$\cdot$	The matrix multiplication operator

Notation: The notation used in this article is summarized in Table I.

II. Literature Review: Interpretability and Performance Gaps

Recent ML-based building control studies offer some of the clearest comparisons between interpretable and closed-box methods. They often start with a high-performing reinforcement learning (RL) or ML controller, then either distill it into an interpretable policy or enhance it with explanation tools. In these cases, the benefits of interpretability frequently outweigh any small performance loss, and some studies even report similar or improved results [28], [29]. For instance, An et al. [30] criticize earlier model-based RL for HVAC because it depends on opaque thermal dynamics models and stochastic optimizers without reliability guarantees. Instead, the authors derive deterministic decision-tree policies from the same models and data. The resulting trees can be verified offline using an HVAC-specific safety criterion and are reported to improve energy savings by 68.4% and comfort by 14.8% compared with the model-based RL baseline. In this case, interpretability supports formal verification and human-readable logic while preserving, and reportedly improving, both energy and comfort performance. More often, interpretable policies perform nearly as well as closed-box baselines while offering practical benefits. Cho and Park [31] develop a deep Q-learning controller for a parallel cooling system, then simplifies its policy into a compact decision-tree-based rule controller. This reduced

controller is easier for facility managers to understand and adjust, with a reported energy-saving difference of about 1.2% compared with the original deep Q-network.

Similarly, Stieger et al. [32] analyze the RL-based HVAC controllers using feature attribution and rule extraction to identify influential sensor inputs and derive rule-based control logic compatible with the existing building automation systems (BASs). The study reports a sensor-count reduction of more than 70% without loss of control performance, showing that interpretability can reduce hardware, commissioning, and maintenance costs while keeping the control logic easier to audit and update under sensor drift or changing occupancy patterns.

Policy distillation into simple schedules further highlights system-level interpretability benefits. Goldfeder and Sipple [33] present a high-performing neural-network RL policy for commercial buildings and distill it into a fixed setpoint schedule that technicians can easily read, adjust, and deploy. Other explainability-focused works keep closed-box controllers but improve transparency. For instance, Kotevska et al. [34] apply interpretability tools to a real-house DRL HVAC controller and report expert validation that the agent's actions match engineering intuition. Similarly, Zhang and Chen [35] combine Shapley values with a large language model to produce narrative explanations of MPC-based precooling decisions consistent with rule-based reasoning. These studies suggest that, in many practical building-control scenarios, the benefits of interpretability can outweigh concerns about a small performance gap. Carefully derived interpretable controllers, such as decision trees, rule sets, or schedules, often achieve energy and comfort performance similar to closed-box RL (often within 0%–1% in reported studies) and, in at least one notable case, outperform a closed-box baseline [30].

For safety-critical and long-lived building applications, the lower validation and maintenance risk of interpretable controllers may outweigh marginal gains from slightly better but opaque policies. More broadly, closed-box controllers can perform well, but their deployment cost increases with sensor needs, tuning effort, and verification challenges. The key question is therefore which modeling choices best balance accuracy, interpretability, and deployability. From this perspective, suitable models should 1) expose a low-dimensional, physically meaningful state and 2) approximate nonlinear dynamics through simple, inspectable local models. dPCA addresses the first requirement by extending standard PCA beyond static variance compression and constructing latent variables that capture dominant temporal modes in multivariate time series [36]. For HVAC and building operation, where thermal dynamics, occupancy, and weather disturbances are important, dPCA provides an interpretable reduced state tailored to control-relevant dynamics. On the regression side, LWR and other lazy learning methods offer accurate and interpretable local alternatives to global closed-box models.

For example, Atkeson et al. [37] show that LWR can model complex nonlinear dynamics through local linear models around the current operating point, while remaining transparent and analyzable. A closely related body of work in process-industry soft sensors also supports a dPCA + LWR architecture, showing that PCA-like latent variables combined with LWR often outperform standard PCA regression as well as purely global or local models.

In spectroscopy, the modified locally weighted PLS methods that use PCA score-space sample selection reduce prediction error by 5%–30% compared with global PLS, highlighting the effectiveness of combining PCA-based latent variables with local regression for nonlinear, time-varying systems [38]. Conceptually, these studies show that: 1) performing local modeling in a low-dimensional latent space is superior to operating directly in the raw high-dimensional space and 2) LWR within that space recovers nonlinear behavior without sacrificing interpretability of the local models. Combining dPCA with LWR for building/HVAC control can be framed as a principled synthesis of the two strands: dPCA provides a compact, dynamically meaningful latent state tailored to building processes, while LWR supplies flexible local linear models in that latent space.

To the best of our knowledge, no building-control study has yet combined dPCA with LWR; most interpretable HVAC controllers instead rely on physics-based gray-box models. Supported by evidence that dPCA outperforms static PCA in dynamic multivariate systems and that PCA-plus-local-regression hybrids work well in related industrial applications, dPCA + LWR can be positioned as a strong compromise between interpretability and performance. It preserves much of the nonlinear modeling capability of closed-box ML while providing a transparent dynamic state and locally linear control laws that are easier to validate, debug, and integrate into real BASs.

The main contributions of this work are summarized as follows.

1) *Constraint-Aware Optimization in Latent (Score) Space for Actionable Control:* Rather than optimizing directly in the original high-dimensional sensor space, we formulate indoor-environment adjustment as a constrained nonlinear optimization in the reduced score space and recover feasible variable shifts that satisfy physical/operational bounds (including the option to freeze nonactuated variables), solved using SQP.
2) *Dynamic Latent Representation With Principled Lag Selection (dPCA):* We extend static PCA to a dynamic formulation via lag augmentation and select the lag window using an SVD-based redundancy criterion, yielding a compact dynamic representation ($l = 4$, i.e., $\approx$1 h at 20-min sampling) suitable for streaming deployment.
3) *Interpretable Local Nonlinear Prediction Under Small-Data Constraints (LWR in Score Space):* We employ LWR in the latent score space to capture nonlinear, context-dependent behavior while preserving interpretability through dPCA loadings and neighborhood influence, avoiding computationally heavy training/hyperparameter tuning typical of closed-box ML/DL models.

The novelty is not a new PCA, dPCA, LWR, or SQP algorithm, but their integration into a constraint-aware score-space

TABLE II
VARIABLES IN **X** AND THEIR UNITS

Variable	unit	Variable	unit
Temp	[°C]	P	[bar]
RH	[%]	CO_2	[ppm]
IAQ	[]	VOC	[ppm]
External Temp	[°C]	External RH	[%]
People density	$[\frac{\#\text{ people}}{\text{m}^2}]$	Static IAQ	[]

optimization framework that maps optimized latent-space corrections back to physically feasible environmental variables.

III. System Architecture

This work relies on data collected using the Tauanito INAIR sensor, a wireless IoT-enabled device developed by Taua for indoor air quality (IAQ) monitoring [26]. The device supports modular sensor integration, includes a built-in GPS for geolocation, and transmits data via cellular networks.

The system uses a Bosch BME680 sensor to monitor gas, RH, pressure (P), and Temp. Volatile organic compounds (VOCs), emitted by industrial and household sources, can impact health and air quality. The sensor also computes an IAQ index ranging from 0 to 500.

The core of the system is an ultralow-power microcontroller unit (MCU), powered by rechargeable lithium-ion batteries. Power consumption is minimized through adaptive algorithms that manage sensor activity and MCU power modes. The system enters ultralow-power states when idle, disabling high-frequency clocks and nonessential components. Other power optimization policies are implemented, including pauses in data transmissions in case of low batteries, local data acquisitions, and automatic device shutdowns. The system includes a GPS module and implements a cellular connection using general packet radio service (GPRS). Data are periodically uploaded to the cloud using the MQTT protocol over transmission control protocol (TCP)/Internet protocol (IP), ensuring a reliable transmission. A cloud-based infrastructure manages the data, offering access through a secure, GDPR-compliant web application with encryption and user authentication.

IV. Background

This section recalls the basic ideas of the proposed framework of the optimization process. Also, refer to [26] for deeper insight into the model on which the optimization is based. The main notations used in this article are summarized in Nomenclature.

A. Variables, Dataset, and Dimensionality Reduction

The model we developed in [26] is based on acquired (by a population poll) wellness observations (satisfaction and PMV) y_i, with $i = 1, \ldots, N$, together with acquired (by an IoT network) environmental variables x_{ij}, where $j = 1, \ldots, K$. In our context, $K = 10$ including VOCs, carbon dioxide (CO_2) and P, and other variables listed in Table II, and $N = 765$ is the number of observation samples in different contexts and sampling times. Therefore, the model aims to predict wellness

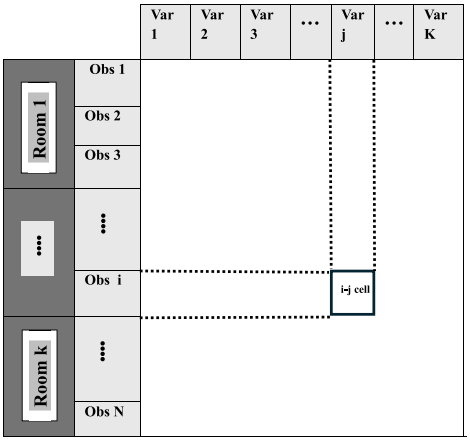

Fig. 3. Example of the dataset matrix. Each cell (i, j) corresponds to variable (Var) j at observation (Obs) i.

$\hat{y}_p$ from new sets of P samples by a nonlinear model $\mathcal{F}$, such that for a generic p sample, it is $\hat{y}_p = \mathcal{F}(x_{p1}, \ldots, x_{pk})$, where $p = 1, \ldots, P$. As wellness metrics, we consider: 1) the PMV of Fanger's model [17], which expresses as a thermal sensation vote (TSV), which is a fundamental tool in the evaluation of thermal comfort, frequently compared with the PMV [39], on a 7-point scale from +3 (hot) to −3 (cold), with 0 denoting neutrality and 2) a self-rated satisfaction score on a 1–10 scale, refer to Appendix A. Both metrics were collected via a questionnaire automatically distributed to a group of users through messaging applications, ensuring minimal intrusion and improving the likelihood of higher response rates. The test involved 46 participants in R_A, 27 in R_B, and 16 in R_C, with the data analyzed in aggregate form. Since the environmental measurements and the control actions are defined at room level, the subjective comfort target was also constructed at the same resolution. The questionnaire-derived satisfaction data were represented as one room-level satisfaction label associated with each environmental observation, rather than as separate occupant-specific labels. This is consistent with the objective of predicting and optimizing the overall perceived comfort of each indoor zone. Because the feedback was sparse and asynchronous, a classical interrater reliability coefficient was not computed. Residual disagreement among occupants is treated as subjective label noise, partially mitigated by the room-level label construction and by the local averaging behavior of LWR.

To reduce temporal leakage during model selection and sensitivity analysis, blocked cross-validation (CV) was used with training-only preprocessing. In each fold, scaling, dPCA projection, and LWR fitting were estimated from the training data only and then applied to the validation block. The initial random 80/20 split is therefore not interpreted as a strictly chronological holdout.

The resulting dataset matrix is denoted as $\mathbf{X} \in \mathbb{R}^{N \times K}$, where N represents the number of time-stamped observations and K corresponds to the total number of recorded variables. A schematic representation of this structured dataset is illustrated in Fig. 3.

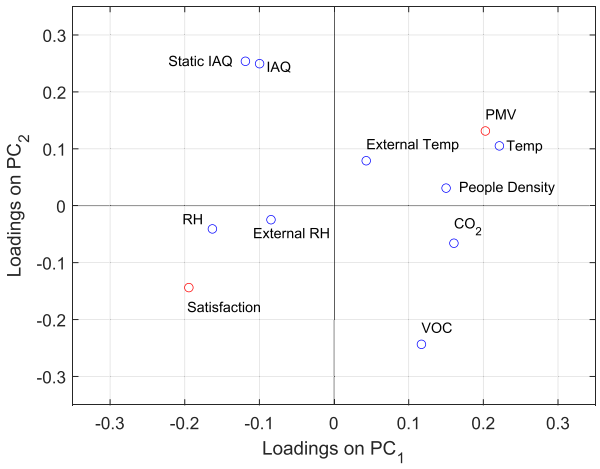

Fig. 4. Loading plot of the first two PCs. Each point is the centroid of all lagged loadings associated with the same variable.

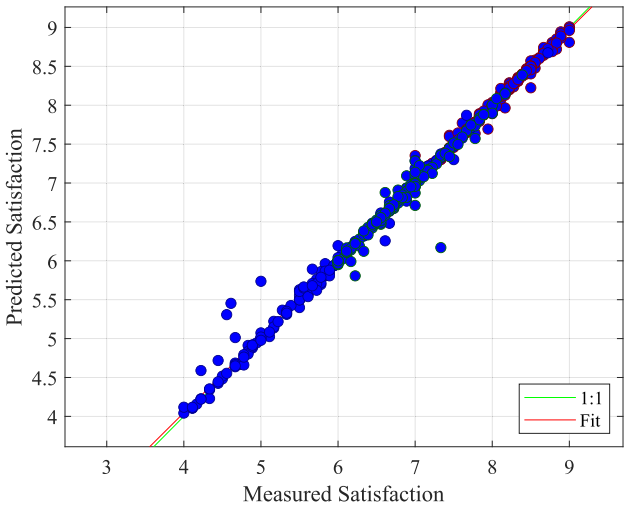

Fig. 5. Predicted versus observed satisfaction in the LWR model [26].

The precalibrated MVA models, such as PCA, are used to identify correlations and reduce dimensionality [40]. The transformation is expressed as a reduced-rank dataset matrix

$$\hat{\mathbf{X}} = \mathbf{T} \cdot \mathbf{L}^{\top} \tag{1}$$

where $\hat{\mathbf{X}} \in \mathbb{R}^{N\times K}$, $\mathbf{L} \in \mathbb{R}^{K\times A}$ is the loadings matrix and $\mathbf{T} \in \mathbb{R}^{N\times A}$ is the score matrix. A is the number of PCs of the dimensionality reduction. Variables that cluster in a so-called "loading plot" are positively correlated, while orthogonal ones are uncorrelated. Therefore, dimensionality reduction allows for decoding hidden properties between variables that could help in optimizing the performance. We will discuss the loading plot of our context later, as shown in Fig. 4. To avoid overfitting, CV is performed [41], ensuring the model generalizes to unseen data. This step is critical for time-series data, which often violates PCA's assumption of independence due to high sampling frequency or inherited autocorrelation [42]. Dimensionality reduction significantly reduces the complexity of the problem, forming the foundation for all optimization structures discussed later. For simplicity of notation, in the following expressions, when we will use the symbol $\mathbf{X}$, we intend a reduced-rank one.

B. Baseline Predictive Model

Before optimization, a model that relates wellness to the ambient variables should be devised to be used as a reference. For this purpose, we used a nonlinear model that we developed in [26] by using LWR, which is here briefly summarized. LWR is a supervised, nonparametric, PCA-based technique used for regression in the presence of nonlinear relationships. Unlike global models, LWR does not use a single set of regression coefficients [37], [43]. Instead, it computes them locally for each prediction point by considering only a subset of neighboring observations. Therefore, rather than using the full dataset, LWR selects the q closest points, determined by a predefined hyperparameter, relative to the target observation. These points are chosen based on a distance metric in the score space. The commonly used metric is the Mahalanobis distance that accounts for the correlation among variables, making it more robust than Euclidean distance in multivariate contexts [44].

In practice, once the q-closest observations are selected and weighted, the regression parameters are computed locally for that region. This piecewise approach makes LWR particularly effective for capturing nonlinear behavior. LWR is a lazy learning method, and the model is built at runtime using neighbor data and kernel weighting, which makes it computationally intensive. However, considering dimensionality reduction (PCA) and limited neighbor sets, LWR could be feasible to be implemented in low-power edge computing techniques [45], [46], [47]. In practice, we selected the LWR hyperparameters by a grid search on the training set only (blocked K-fold CV), varying the neighborhood size q and latent dimension A. The performance was stable over a broad region around the selected operating point $(A, q) = (7, 35)$, and the Mahalanobis distances exhibited similar trends, indicating low sensitivity to the choice of distance in the score space.

In the adopted LWR implementation, the neighbor weights are not learned as additional model parameters. They are computed using a fixed tri-cube kernel after selecting the q nearest neighbors in the reduced score space. For each query point, the bandwidth is adaptive and corresponds to $h = d_q$, i.e., the distance to the qth nearest neighbor [23]. Thus, the kernel shape is fixed, while the main bandwidth-related parameter is the neighborhood size q. The bandwidth sensitivity was evaluated using the unified pooled model and reported across rooms R_A, R_B, and R_C. Very small neighborhoods produced unstable errors, whereas $q = 35$ lay in the stable moderate-neighborhood region for the three room contexts. Since larger q values increase the adaptive bandwidth and reduce locality, $q = 35$ was retained as a common accuracy–locality compromise.

To further assess real-time feasibility and possible failure modes, we evaluated LWR at the selected operating point $(A, q) = (7, 35)$. A test point was considered out-of-distribution (OOD) when its distance to the qth nearest training neighbor exceeded the 95th percentile of the corresponding training-set distance distribution.

TABLE III
LWR REAL-TIME FEASIBILITY AND FAILURE-MODE DIAGNOSTICS AT THE SELECTED OPERATING POINT $(A, q) = (7, 35)$

Metric	RMSECV	RMSEP	Time(s)/samp.	OOD	Eff.
Euclidean	0.4183	0.2925	6.47×10^{-5}	4.05%	13.70
Mahalanobis	0.3743	0.2436	6.90×10^{-5}	4.05%	12.81

Table III shows that the prediction time remains on the order of 10^{-5} s/sample, confirming real-time feasibility after projection into the reduced score space. The OOD ratio is only 4.05%, the effective neighborhood size (Eff) remains around 13 samples, and no ill-conditioned local fits are detected. The Mahalanobis metric also gives lower RMSECV and RMSEP than Euclidean distance, supporting its use in the final implementation.

Fig. 5 shows the regression curve generated by the LWR model of the actual versus the predicted satisfaction metric [26]. This insight provides a deeper understanding of the local context of the observations and guides the weighting scheme applied in the LWR analysis. Therefore, the optimization technique will be aimed at moving points from lower wellness levels to higher ones.

V. DPCA IMPROVEMENTS

The time series of environmental variable samples could exhibit short-term memory dependency. In this section, we will explore improvements in variable reduction using the dPCA technique.

A. dPCA Construction

The object of this technique is to incorporate not only the current values of environmental parameters but also several previous time samples. This structure enables the capture of temporal dependencies critical to indoor environmental dynamics. It was first introduced by Brillinger [48] and further refined in subsequent studies [36]. Traditional PCA struggles with identifying specific relationships in time-dependent data due to its static nature. dPCA addresses this by combining both autocorrelation and cross correlation, enabling a more comprehensive analysis of multivariate time series [42]. The core idea of dPCA involves augmenting the original dataset by including lagged versions of the input variables, forming a matrix that captures both current and past states

$$\mathbf{X}^{(l)}(t) = [\mathbf{X}(t), \mathbf{X}(t-1), \ldots, \mathbf{X}(t-l)] \tag{2}$$

where l is the generic time lag.

To determine the optimal number of lags, an iterative algorithm is used based on the decomposition of the augmented data matrix via SVD. This is motivated by the observation that low singular values correspond to linear dependencies that limit the variability in the data [49]. The algorithm is here summarized in pseudo-code [42]. The notations and symbols used are presented in Nomenclature.

In practice, due to the presence of noise, singular values rarely reach exactly zero. Therefore, the values below a small threshold *Thr* are treated as indicators of redundant (linear)

Algorithm 1 dPCA Algorithm for Optimal Lag Selection

$l \leftarrow 0$
$r_{\text{new}}(l) \leftarrow 1$
while $r_{\text{new}}(l) > 0$ **do**
  $X_{\text{dynamic}} \leftarrow [X(t), X(t-1), \ldots, X(t-l)]$
  $P \leftarrow \text{PCA}(X_{\text{dynamic}})$
  $S \leftarrow \text{SVD}(P)$
  $r(l) \leftarrow \text{count}(S \le \text{Thr})$
  $l \leftarrow l + 1$
  $r_{\text{new}}(l) \leftarrow r(l) - \sum_{i=0}^{l-1}(l-i+1) \cdot r_{\text{new}}(i)$
end while

structure in the data [42]. In here, **P** denotes the PC matrix derived from applying PCA to the augmented dataset **X**. The vector $r(l)$ counts how many singular values in **S** fall below the threshold Thr at lag l, indicating linear constraints. The residual r_{new} measures the marginal gain from including one more lag. The optimal lag is found when $r_{\text{new}} > 0$, meaning no further linear dependency is being added.

In this study, the algorithm determined an optimal lag of $l = 4$, which corresponds to a time window of 1 h. The augmented matrix can be used with various MVA techniques, such as PCA and LWR. PCA is employed to reduce dimensionality by identifying the directions of maximum variance. LWR is used specifically for nonlinear regression, predicting satisfaction scores based on the locally weighted influence of past observations and was used here as a baseline predictive model. These dynamic methods are particularly suitable in environments, where variables interact over time. In Section VI these approaches are applied to predict the optimal Temp that maximizes occupant satisfaction. All datasets are autoscaled prior to analysis.

Although the spectral analysis of the environmental variables revealed a strong daily periodicity (≃72 lags at 20 min sampling), the optimal lag selected was $l = 4$ (≃1 h). This is expected, since the singular-value-based criterion in dPCA seeks the shortest lag window that captures incremental autocorrelation and cross correlation without introducing linear redundancy. Longer lags, such as those corresponding to circadian cycles, are partly predictable from shorter ones and thus fail to produce new low-singular-value directions in the augmented matrix. To verify that this choice is not an artifact of a specific sampling rate or room, we repeated the lag sweep after downsampling the streams (e.g., $\Delta t \in \{40, 60\}$ min) and by evaluating the rooms separately (R_A, R_B, and R_C). In all cases, root mean square error for CV (RMSECV) decreased up to a moderate memory length and then saturated, indicating diminishing returns for larger l due to redundancy in the lag-augmented space; hence, $l = 4$ was retained as the most robust accuracy–complexity compromise.

Table IV reports the RMSECV values at representative lag values $L = 2$, $L = 4$, and $L = 6$ for each temporal resolution and room context. The comparison shows that the native temporal resolution, $ds = 1$ (20 min), generally provides the lowest RMSECV values, confirming that downsampling can attenuate short-term information relevant for satisfaction prediction. At the native resolution, $L = 4$ is the best among

TABLE IV
RMSECV VALUES FOR SELECTED LAG VALUES $L = \{2, 4, 6\}$ UNDER DIFFERENT TEMPORAL RESOLUTIONS AND ROOM CONTEXTS

Context	ds=1 (20 min)			ds=2 (40 min)			ds=3 (60 min)		
	$L=2$	$L=4$	$L=6$	$L=2$	$L=4$	$L=6$	$L=2$	$L=4$	$L=6$
All rooms	0.3161	0.3481	0.3949	0.5629	0.5676	0.5224	0.5723	0.7051	0.6072
Room R_A	0.2162	0.1933	0.3212	0.3403	0.3422	0.3498	0.3546	0.3496	0.3594
Room R_B	0.6525	0.5098	0.6097	0.7868	0.6788	0.5665	0.7010	0.9003	0.8260
Room R_C	0.2868	0.2063	0.1828	0.2800	0.3316	0.3902	0.4131	0.4311	0.3702

the three tested representative lags for rooms R_A and R_B and remains in the low-error region for the aggregate dataset and room R_C. Although $L = 6$ gives a slightly lower RMSECV for R_C, it increases the dimension of the lag-augmented space. Therefore, $L = 4$ is retained as a robust accuracy–complexity compromise: it captures useful short-term dynamics while limiting the dimensional growth associated with longer lag windows. For deployment, the lag was selected using a unified accuracy–complexity rule rather than tuned independently for each room. Specifically, among the representative lag values tested at the native sampling resolution $ds = 1$, $L = 4$ provided the best average RMSECV across the roomwise models while limiting the dimensional growth of the lag-augmented matrix. The same lag value $L = 4$ was therefore applied uniformly across all rooms to maintain a consistent and generalizable deployment configuration.

To assess whether the learned dynamic latent structure reflects different occupancy regimes, we fitted the same dPCA configuration ($l = 4$ and $A = 7$) separately on each room subset. The resulting variance spectra differ across rooms: in the higher occupancy spaces (R_A and R_B), variance is more concentrated in PC1, whereas the staff room (R_C) exhibits a less concentrated spectrum with more variance distributed over higher PCs. To connect these differences to occupancy, we inspected the mean absolute loading of the people density channel across components. Occupancy-related variability is captured by different dynamic latent modes depending on room usage (e.g., a dominant people density contribution around PC3 for R_B and stronger contributions in higher PCs for R_A). These observations support the use of dPCA: by embedding short-term temporal structure, dPCA yields a compact dynamic state whose dominant modes adapt to room-specific occupancy dynamics, rather than relying on a single static variance structure across all zones. As a result, the algorithm prioritizes compact short-term dynamics over longer term seasonal patterns. Longer term periodicity could be included with some techniques in dPCA, but it has not been addressed in this article.

B. Effectiveness of dPCA

Both static and dynamic modeling approaches were applied to the same dataset, and their performance was evaluated using three standard metrics: root mean square error for calibration (RMSEC), RMSECV, and root mean square error for prediction (RMSEP). The results are summarized in Table V.

TABLE V
PERFORMANCE METRICS OF THE STATIC AND DYNAMIC METHODS

		#PC	RMSEC	RMSECV	RMSEP
Static	PCA	2	0.7933	0.8572	–
	LWR	2	0.8348	0.6860	0.7798
Dynamic	PCA	7	0.2572	0.3950	–
	LWR	7	0.1305	0.3084	0.3570

To evaluate the effectiveness of dPCA, CV was performed using a calibration set comprising 80% of $\mathbf{X}_{\text{dynamic}}$ and a test set comprising the remaining 20%, selected using a stratified sampling strategy.

The results indicate that the dynamic approach consistently outperforms the static one in both the decomposition and nonlinear regression tasks. In particular, dynamic LWR achieves the lowest error across all metrics, with an RMSEC of 0.1305, RMSECV of 0.3084, and RMSEP of 0.3570, compared to its static counterpart (RMSEC = 0.8348, RMSECV = 0.686, and RMSEP = 0.7798). This demonstrates the advantage of modeling variables as time-dependent processes, which allows dynamic methods to better capture temporal patterns and lagged dependencies. These findings support the hypothesis that the satisfaction score exhibits nonlinear behavior with respect to environmental variables. For example, extremely low or high Temp can both lead to dissatisfaction, resulting in a nonmonotonic relationship that cannot be effectively captured by linear models [26]. Dynamic models, by including lagged inputs, are more capable of addressing these complexities.

To isolate temporal augmentation from spatial correlation reduction, we evaluated four variants [raw LWR, PCA(score) + LWR, lagged-LWR, and dPCA(score) + LWR] under identical splits. The best performance is achieved by *dPCA(score)+LWR* ($L = 4$), with RMSEP = 0.1459 (R_C), 0.5093 (R_A), and 0.3794 (R_B), outperforming both *PCA(score)+LWR* ($L = 0$) and *raw LWR* ($L = 0$). In contrast, *lagged-LWR* without PCA compression degrades performance (e.g., R_B: 0.6837 to 0.8244), confirming that the gain comes from combining temporal memory with score-space compression rather than from raw lag concatenation.

This ablation confirms that the improvement is not obtained by simply adding lagged variables. Raw temporal augmentation without score-space compression increases dimensionality and degrades the local regression, whereas the combination of lag augmentation and PCA-based compression provides a compact dynamic representation that improves prediction

accuracy. Therefore, the benefit of dPCA comes from the joint effect of temporal memory and spatial correlation reduction.

Further insight is provided by the loading plot in Fig. 4, generated through dPCA, which illustrates the correlation between variables. Variables positioned close together in the plot are positively correlated, whereas those on opposite sides exhibit negative correlation. Notably, PMV and Temp are closely aligned, indicating a strong positive correlation, while satisfaction scores show a negative correlation with both, further validating the modeling assumptions.

The choice of LWR is supported by a direct comparison with alternative nonlinear regressors on the same score-space inputs, showing a favorable accuracy–complexity tradeoff in our setting. The Kernel regression (Nadaraya–Watson), Gaussian process regression (GPR), and gradient-boosted decision trees (GBDT) were evaluated on the same score-space inputs and preprocessing as LWR (training-only standardization; $l \in \{0, 4\}$; $A = 2$ static and $A = 7$ dynamic), using blocked-fold RMSECV and a chronological holdout RMSEP. LWR achieves the lowest errors (static: RMSECV $\approx$ 0.686, RMSEP $\approx$ 0.780; dynamic: RMSECV $\approx$ 0.308, and RMSEP $\approx$ 0.357), while the alternative nonlinear models yield higher errors; dynamic augmentation improves all methods. Recent attention-based saliency models provide a useful perspective for interpreting input relevance in prediction models. For example, Khan et al. [50] proposed an attention-enhanced saliency-detection framework that emphasizes the most informative visual regions and feature channels through attention mechanisms. While that work is developed in the image domain, its central principle of assigning higher importance to the most relevant inputs is also meaningful for multivariate sensor-based prediction. In the proposed framework, this relevance mechanism appears through two complementary levels: dPCA loadings indicate which environmental variables contribute most to the latent dynamic representation, while LWR distance-based neighborhood weights indicate which historical operating points most influence each predicted satisfaction value. Thus, the dPCA+LWR formulation provides an interpretable and lightweight way to identify relevant variables and local operating regimes in IoT sensor streams, following the same general motivation of saliency-based methods: making the prediction process more transparent by highlighting influential inputs.

C. Computational Cost and Comparison With Alternative Temporal Models

To quantify the computational overhead of the proposed dynamic representation, we report the dominant (big-$\mathcal{O}$) time and memory costs for the main method families discussed in this work. Let N denote the number of time samples, k the number of measured variables, let V be the number of candidate ARIMA orders tested with G optimizer iterations, and for recurrent models, let M be the number of training sequences/windows, H the hidden size, d the input dimension, and E the number of training epochs. Under these definitions, static PCA scales as $\mathcal{O}(Nk^2 + k^3)$ (dense covariance/SVD) [51], while lag-augmented dPCA scales as $\mathcal{O}(Nk^2l + (k(l + 1))^3)$, where the second term reflects the conservative dense eigendecomposition/SVD cost in the lag-stacked space [52]. ARIMA identification with model-selection scales as $\mathcal{O}(VGN)$ for a single series (and linearly with the number of series if fitted independently per variable) [53], whereas recurrent neural network (RNN)/long short-term memory (LSTM) training scales as $\mathcal{O}(EMN(H^2 + Hd))$, reflecting iterative training and backpropagation through time [54]. In our setting ($N = 765$, $k = 10$, and $l = 4$), the dPCA overhead therefore amounts to a constant-factor increase over static PCA, while remaining substantially lighter than recurrent temporal models. The state-space identification is also considered as an additional temporal baseline [55]. Since no calibrated thermal state-space model is available for the monitored rooms, a fair state-space comparison must include the cost of identifying such a model from data. Subspace-type identification relies on block-Hankel construction and SVD-based subspace extraction; therefore, we use the conservative proxy $\mathcal{O}((qk)^2N + (qk)^3 + C_sI_sN(s^3 + ks^2))$, where q is the block-Hankel depth, s is the latent state dimension, and C_s and I_s denote the number of candidate state-space structures and refinement iterations, respectively. Fig. 6 compares the instantiated time proxies at the dataset scale ($N = 765$ and $k = 10$). The multivariate methods (PCA and dPCA) lie in the lowest compute regime, since their cost is dominated by a single (offline) linear-algebra decomposition on a low-dimensional embedding. State-space identification and ARIMA model selection require higher time proxies due to the model-identification/search steps, while recurrent models (RNN/LSTM) exhibit markedly higher time proxies due to iterative training over multiple epochs and the $\mathcal{O}(H^2)$ recurrent-state term. In terms of memory, dPCA remains close to PCA, whereas state-space identification requires storage for block-Hankel/subspace matrices and RNN/LSTM requires substantially higher storage due to backpropagation-through-time activations. Overall, the proposed dPCA-based pipeline remains computationally lightweight relative to deep temporal baselines while still capturing short-term dynamics.

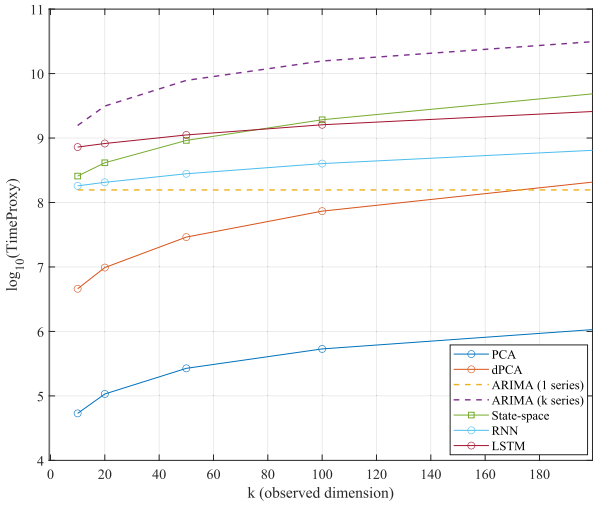

Fig. 6. Time proxy at the dataset scale ($N = 765$, $k = 10$, $A = 7$, $l = 4$, $H = 64$, and $E = 50$).

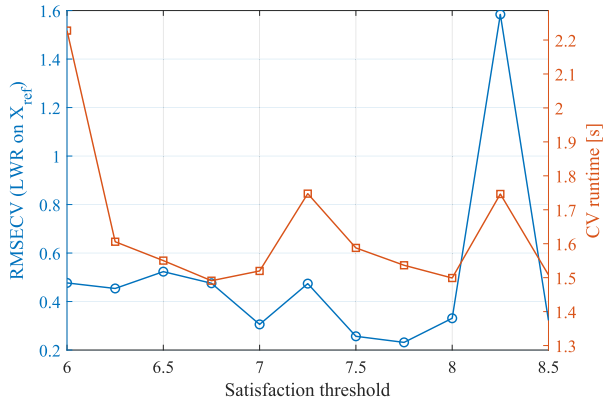

Fig. 7. Grid search for the satisfaction threshold used to define the reference set $\mathbf{X}_{\text{ref}}$. The left axis (blue) reports RMSECV of the LWR model trained on $\mathbf{X}_{\text{ref}}$, while the right axis (orange) reports the corresponding CV runtime.

VI. Wellness Optimization by Environmental Variable Change

A. Score Space Optimization

The optimization approach is based on a previously developed nonlinear LWR model constructed on a reduction of variables where observations are projected on PCs into the score space [26]. We have shown that dPCA further enhances variable reduction. This model enables the prediction of wellness from the environmental variables by computing a weighted sum of the nearest neighboring observations in the score space, with weights defined locally. One of the main features of this approach is to allow a representation in reduced score space, offering great simplification to the optimization calculations.

The optimization algorithm proceeds based on the following steps. First, a subset of observations with the highest wellness values is selected as the target. Then, the out-of-target observations are adjusted, within the predictive model, to move as close as possible to the target group. The changes required in the environmental variables (within predefined constraints) to reach the target indicate the optimal configuration for maximum satisfaction. Since the problem involves a nonlinear function subject to constraints on the variables, it falls within the domain of nonlinear programming and regularization.

Specifically, a dataset $\mathbf{X}$, composed of the environmental variables listed in Table II, is used to identify and model this region, using the satisfaction scores collected from a survey. From $\mathbf{X}$, two subdatasets are constructed: the first dataset, $\mathbf{X}_{\text{ref}}$, includes all observations whose satisfaction scores exceed a predefined threshold. A grid search was conducted to determine the optimal threshold value. Specifically, the grid search in Fig. 7 was performed over thresholds from 6.00 to 8.50 with a step size of 0.25, resulting in 11 candidate values. For each threshold, X_{ref} was rebuilt and the corresponding LWR CV error and runtime were evaluated. The final value, $\tau = 7.5$, was selected as the lowest threshold within the high-satisfaction low-error plateau around 7.5–7.8, rather than as the global minimum over the full grid. This choice preserves a larger reference set, improves local-neighborhood support for LWR, and maintains a favorable accuracy–complexity balance. The second dataset, $\mathbf{X}_{\text{obj}}$, contains the remaining observations. An LWR model is trained using six latent variables on the reference dataset (RMSEC = 0.0641 and RMSEP = 0.1564) to predict satisfaction scores. The model projects both datasets into score space, enabling comparison in a reduced-dimensional setting. A score vector $\mathbf{T}'_{\text{obj}} \in \mathbb{R}^{1\times A}$ is linearly updated when its corresponding observation row $(\mathbf{X}_{\text{obj}}) \in \mathbb{R}^{1\times K}$ changes by a vector $\mathbf{\Delta} \in \mathbb{R}^{1\times K}$.

Algorithm 2 Environmental Adjustment via Score-Space Optimization

model $\leftarrow$ LWR($\mathbf{X}_{\text{ref}}$)
$\mathbf{T}_{\text{ref}} \leftarrow$ mean(model.scores)
$\mathbf{L} \leftarrow$ model.loadings
$\Delta \leftarrow$ zeros(size(x))
for $n = 1$ to rows($\mathbf{X}_{\text{obj}}$) **do**
    $\mathbf{x} \leftarrow \mathbf{X}_{\text{obj}}(n)$
    $\mathbf{T}_{\text{obj}} \leftarrow$ model($\mathbf{x}$)
    $\mathbf{T}'_{\text{obj}} = \mathbf{T}_{\text{obj}} + \mathbf{\Delta} \cdot \mathbf{L}$
    dist $\leftarrow \|\mathbf{T}_{\text{ref}} - \mathbf{T}'_{\text{obj}}\|$
    $\mathbf{\Delta} \leftarrow$ SQP(dist)
    $\mathbf{x}_{\text{new}}(n,:) \leftarrow \mathbf{x}(n,:) + \mathbf{\Delta}(n,:)$
end for

Using nonlinear programming approach [56], [57], the objective is to compute the optimal $\mathbf{\Delta}$ that shifts the score $\mathbf{T}'_{\text{obj}}$ as close as possible to the mean reference score $\mathbf{T}_{\text{ref}} \in \mathbb{R}^{1\times A}$. The optimization problem includes an L_2 regularization term for stability and robustness against outliers [58] and respects physical constraints (e.g., some variables cannot be altered directly). The L_2 regularization weight was selected by an offline sensitivity sweep over a logarithmic grid and by evaluating the tradeoff between 1) improvement toward the reference in score space and 2) the magnitude of the required variable shifts. The resulting Pareto curve exhibits diminishing returns beyond a moderate regularization level; both a knee-point (Kneedle) criterion and an ideal-point criterion selected $\lambda = 1$ as the most balanced operating point, yielding meaningful wellness improvement while avoiding excessively large corrections and preserving stable, feasible SQP solutions across the tested range. The notations of (3) are presented in Nomenclature

$$\begin{aligned}
\mathbf{\Delta} \doteq \ & \arg\min_{\mathbf{\Delta}} \left(\left|\left| \boldsymbol{T}_{ref} - \boldsymbol{T'}_{obj}(\mathbf{\Delta}) \right|\right|_2^2 \right) + \left|\left|\Delta\right|\right|_2^2 \\
\textbf{s.t.} \quad & 0\ ^\circ\text{C} \le Temp + \mathbf{\Delta}_{Temp} \le 30\ ^\circ\text{C} \\
& 0\% \le RH + \mathbf{\Delta}_{RH} \le 100\% \\
& 0 \le IAQ + \mathbf{\Delta}_{IAQ} \le 500 \\
& 0 \le static\ IAQ + \mathbf{\Delta}_{static\ IAQ} \le 500 \\
& CO_2 + \mathbf{\Delta}_{CO_2} \ge 0\ ppm \\
& \mathbf{\Delta}_{VOC} = 0\ ppm \\
& \mathbf{\Delta}_{External\ Temp} = 0\ ^\circ\text{C} \\
& \mathbf{\Delta}_{External\ RH} = 0\% \\
& \mathbf{\Delta}_{P} = 0.
\end{aligned} \tag{3}$$

The IAQ and CO_2 bounds in the SQP are imposed as hard feasibility limits consistent with the sensor domains and

the observed operating range (IAQ/static-IAQ are proprietary indices defined in the software libraries of the adopted Bosch bME280 sensor, on a bounded $[0, 500]$ scale, while CO_2 is constrained to nonnegative and optionally capped to a conservative indoor ventilation-adequacy threshold to avoid unrealistic setpoints). This constrained optimization is solved using an SQP approach. SQP iteratively linearized the objective and constraints, solving a quadratic approximation at each step to find the optimal $\mathbf{\Delta}$. Greater details on the approach could be found in Appendix B. Once obtained, $\mathbf{\Delta}$ is applied to the original observation $\mathbf{X}_{\text{obj}}$, generating the adjusted instance $\mathbf{x}_{\text{new}}$ while maintaining feasibility within physical bounds. The pseudo-code for this optimization routine is shown in Algorithm 2, and the notations and symbols used are presented in Nomenclature.

It is worth noting that the proposed optimization framework can be naturally extended to incorporate energy-awareness. In particular, an additional constraint on the estimated energy consumption $E(\mathbf{x} + \mathbf{\Delta})$ of the HVAC system can be introduced, such that $E(\mathbf{x} + \mathbf{\Delta}) \leq E_{\max}$, where $E_{\max}$ represents a predefined operational budget or peak demand limit. This formulation enables the algorithm to trade off wellness improvement and energy usage, ensuring that comfort optimization is achieved without exceeding desired consumption thresholds.

In Algorithm 2, the update $\mathbf{\Delta} \leftarrow \text{SQP(dist)}$ denotes a call to the numerical SQP solver for the constrained problem in (3). The Lagrange multipliers associated with the active constraints, the active-set changes, and the step-size/line-search parameters are therefore updated internally by the solver. No explicit user-defined dual-variable step size is imposed in the algorithm. Each solve is initialized from a feasible zero-correction point, clipped to the imposed bounds. The numerical convergence diagnostics include the 85.76% success rate and the 15.62/16/19 mean/median/maximum iteration counts for successful solves. The SQP optimization demonstrated robust behavior in our experiments. All returned solutions satisfied the prescribed bounds (zero maximum and mean bound violation), indicating reliable constraint handling. Convergence was stable, with successful solves typically requiring about 16 iterations (maximum 19) and an overall solver success rate of 85.76%. In terms of effectiveness, the median score-space distance to the reference decreased from 5.142 to 4.066 (20.93% improvement), and 68.61% of the objective samples moved closer to the reference. Finally, the postoptimization acceptance rate reached 87.97% on the full dataset and 70.23% on the objective subset.

B. Results

This section evaluates the impact of the proposed algorithm on indoor thermal comfort, focusing on score-space transformation, Temp optimization, and predicted satisfaction levels.

Fig. 8 provides a visualization of the score-space transformation using, for simplicity, a dimensionality reduction to 2 PCs. Purple points indicate high-satisfaction observations used to compute the reference point, while blue points represent low-satisfaction observations targeted for improvement. After the algorithm is applied, the blue points shift toward the purple cluster, confirming the effectiveness of the optimization in aligning target observations with the desired comfort zone.

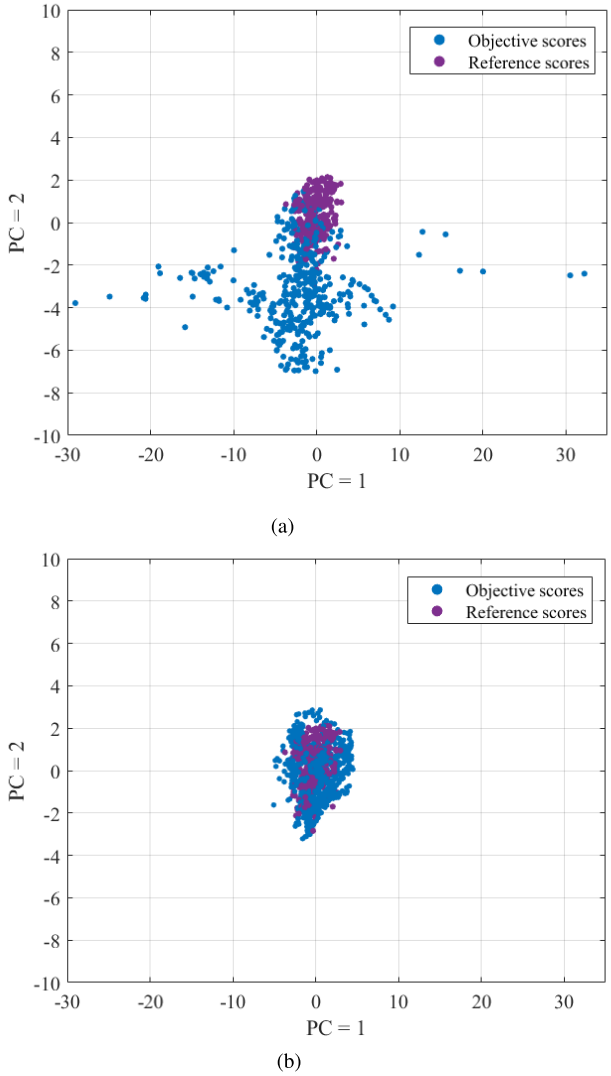

Fig. 8. Score plots of points before and after the optimization. (a) Before optimization. (b) After optimization.

The results of the optimal Temp prediction using the proposed algorithm are shown in Fig. 9(a). In practice, the figure shows the suggested variables to maximize the wellness within all the given constraints. A statistical comparison between the predicted and actual Temp was conducted for each room. In room R_C, the difference was not statistically significant ($p_C = 0.18$), indicating strong agreement between the predicted and actual values. This result is expected, as R_C already had relatively high satisfaction levels prior to optimization and therefore required only minimal correction. Conversely, statistically significant differences were observed in R_A and R_B. These substantial deviations demonstrate the algorithm's ability to shift environmental conditions toward optimal comfort, particularly in spaces initially exhibiting low satisfaction. Fig. 9(b) illustrates the actual versus the predicted optimal RH, showing that the optimized values converge toward the comfort range associated with maximum satisfaction, while remaining within recommended indoor humidity limits.

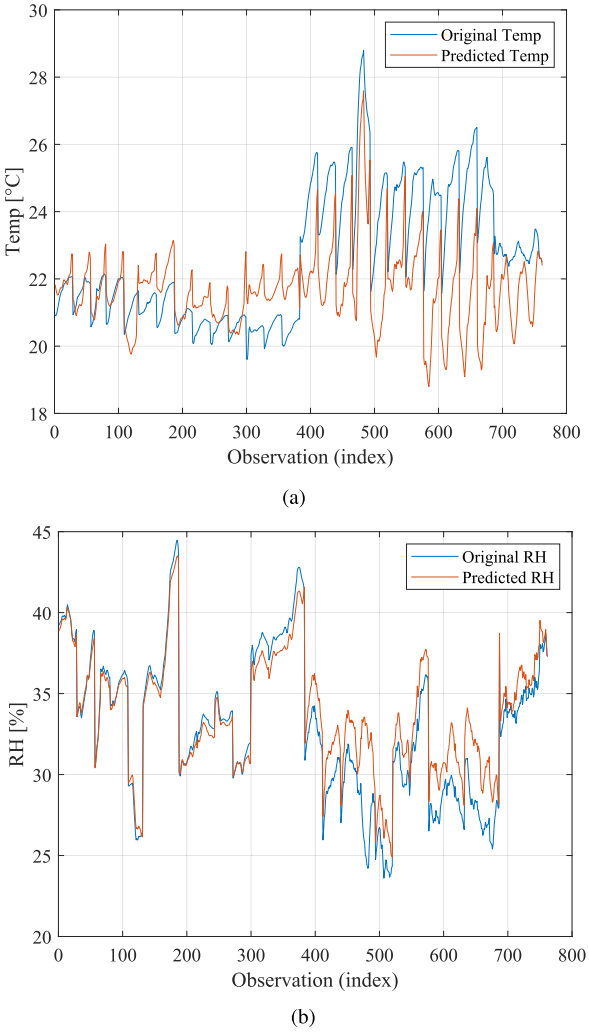

Fig. 9. Prediction of optimal Temp and RH using the proposed optimization compared to the original measured variable values. (a) Optimal Temp. (b) Optimal RH.

The model also affects other environmental variables, such as IAQ indicators, including CO_2 and VOCs concentrations. These parameters are not considered fixed; rather, they evolve within their respective boundary conditions as part of the optimization process. In practice, IAQ (CO_2, VOCs) could be regulated primarily through ventilation strategies, while RH could be controlled via humidification systems within the HVAC.

Nevertheless, Temp is chosen as the primary focus of this study because it is the most influential parameter in thermal comfort perception and directly linked to the PMV index, which serves as our main wellness metric. The impact of the optimized Temp on occupant satisfaction is further assessed through the distribution of PMV scores in Fanger's model, as illustrated in Fig. 10. For insights into Fanger's model, refer to Appendix A.

The predicted PMV values exhibit a mean of -0.27 with a standard deviation of 0.22. Notably, 85.51% of the samples fall within the thermal comfort zone (defined as PMV within ± 0.5), and the dispersion is moderate relative to this band:

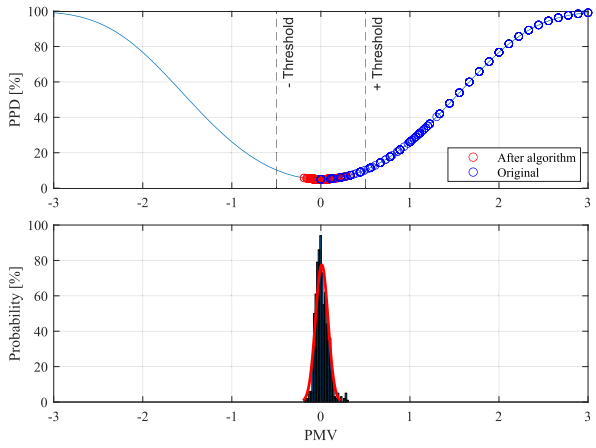

Fig. 10. Distribution of predicted percentage of dissatisfied (PPD) values derived from the predicted PMV scores using Fanger's model.

TABLE VI

SATISFACTION PERCENTAGES FOR EACH ROOM, BEFORE AND AFTER APPLYING THE ALGORITHM

Room	Before	After
R_A	5.18%	94.81%
R_B	12.69%	90.47%
R_C	58.48%	72.84%

$2\sigma \approx 0.44 < 0.5$, indicating that the optimized distribution is tightly concentrated around the comfort threshold. The remaining out-of-band cases are mainly attributable to environmental variability (e.g., outdoor and occupancy-driven disturbances) and to the imposed feasibility/actuation constraints that prevent aggressive corrections, highlighting a substantial improvement in predicted thermal satisfaction across all rooms. A detailed breakdown of satisfaction percentages, both before and after applying the algorithm, is provided in Table VI. Room R_A showed the most significant improvement, increasing from 5.18% to 94.81%. Rooms R_B and R_C also experienced marked gains, rising from 12.69% to 90.47% and from 58.48% to 72.84%, respectively. All reported postalgorithm satisfaction improvements are counterfactual model-based estimates obtained by re-evaluating the trained predictor on the optimized trajectory $\mathbf{X}_{\text{new}}$; they do not represent measured postintervention occupant feedback because no closed-loop HVAC actuation was performed during data collection.

VII. OFFLINE PROXY ANALYSIS OF THERMAL DEMAND AND COMFORT COMPLIANCE

In the absence of metered HVAC energy, we report an offline proxy-based energy analysis using degree-hour indices derived from indoor/outdoor trajectories [59], [60], [61], [62], [63]. The purpose is not to claim kWh savings, but to quantify whether the counterfactual optimized trajectory reduces sensible and latent load proxies. We compare two time-aligned trajectories over the same horizon of $N = 765$ samples with sampling time $\Delta t = 20$ min (253.67 h): baseline measured conditions (superscript (0)) and the counterfactual optimized conditions produced by Algorithm 2 (superscript (1)). First, we define the total sensible-load proxy.

TABLE VII
OFFLINE PROXY SUMMARY (BASELINE VERSUS COUNTERFACTUAL ALGORITHM 2). SAVINGS ARE PROXY REDUCTIONS (NOT METERED KWH)

Metric	Baseline (0)	Algorithm 2 (1)
L_{tot} (°C·h)	3833.83	3705.00
Saving_L (%)	–	**3.36**
$J_{\Delta\omega}$ (kg/kg·h)	1.2402×10^{-1}	1.0819×10^{-1}
Saving_ω (%)	–	**12.76**
D_T (°C·h)	138.64	112.28
H_T (h)	130.00	154.67
$\bar{\delta}_T$ (°C)	1.066	0.726
D_{RH} (%·h)	179.90	125.63
H_{RH} (h)	74.33	65.33
$\bar{\delta}_{RH}$ (%)	2.420	1.923
$H^+_{CO_2}$ (h)	0.00	0.00
$J_{\Delta T}$ (°C·h)	–	202.14
$J_{\Delta RH}$ (%·h)	–	119.47

Sensible-load proxy (temperature difference): Let $T^{(s)}_{in,k}$ and $T_{out,k}$ denote indoor and outdoor temperatures at time k for scenario $s \in \{0,1\}$. We define heating- and cooling-like degree-hour proxies

$$L^{(s)}_{\mathrm{heat}} = \sum_{k=1}^{N} \max\left(0,\, T^{(s)}_{in,k} - T_{out,k}\right) \Delta t \tag{4}$$

$$L^{(s)}_{\mathrm{cool}} = \sum_{k=1}^{N} \max\left(0,\, T_{out,k} - T^{(s)}_{in,k}\right) \Delta t \tag{5}$$

$$L^{(s)}_{\mathrm{tot}} = L^{(s)}_{\mathrm{heat}} + L^{(s)}_{\mathrm{cool}} = \sum_{k=1}^{N} \left|T^{(s)}_{in,k} - T_{out,k}\right| \Delta t. \tag{6}$$

These indices are not energy (kWh), but they are proportional to the dominant sensible load term under unchanged envelope/ventilation coefficients. We report the relative change

$$\mathrm{Saving}_L\,[\%] = 100\left(1 - \frac{L^{(1)}_{\mathrm{tot}}}{L^{(0)}_{\mathrm{tot}}}\right) \tag{7}$$

we obtain $L^{(0)}_{\mathrm{tot}} = 3833.83\ ^\circ\mathrm{C}\cdot h$ and $L^{(1)}_{\mathrm{tot}} = 3705.00\ ^\circ\mathrm{C}\cdot h$, corresponding to $\mathrm{Saving}_L = 3.36\%$ (Table VII).

Latent/ventilation-oriented proxy (humidity ratio): From indoor Temp $T^{(s)}_{in,k}$, indoor $\mathrm{RH}^{(s)}_k$, and outdoor $\mathrm{RH}_{out,k}$, we compute humidity ratio ω and define

$$J^{(s)}_{\Delta\omega} = \sum_{k=1}^{N} \left|\omega^{(s)}_{in,k} - \omega_{out,k}\right| \Delta t \tag{8}$$

$$\mathrm{Saving}_\omega\,[\%] = 100\left(1 - \frac{J^{(1)}_{\Delta\omega}}{J^{(0)}_{\Delta\omega}}\right). \tag{9}$$

This is a latent-load proxy (kg/kg·h), not metered humidification energy. In our data, $J^{(0)}_{\Delta\omega} = 1.2402 \times 10^{-1}$ and $J^{(1)}_{\Delta\omega} = 1.0819 \times 10^{-1}$ kg/kg·h, yielding $\mathrm{Saving}_\omega = 12.76\%$ (Table VII).

Constraint compliance (comfort and IAQ): Given comfort bands $[T_{\min}, T_{\max}]$ and $[\mathrm{RH}_{\min}, \mathrm{RH}_{\max}]$ and an upper limit $CO_{2,\max}$, we compute 1) violation degree-hours and 2) violation hours

$$D^{(s)}_T = \sum_{k=1}^{N}\Big[\max\left(0, T^{(s)}_{in,k} - T_{\max}\right) + \max\left(0, T_{\min} - T^{(s)}_{in,k}\right)\Big] \Delta t \tag{10}$$

$$H^{(s)}_T = \sum_{k=1}^{N} \mathbb{I}\left\{T^{(s)}_{in,k} < T_{\min} \ \vee \ T^{(s)}_{in,k} > T_{\max}\right\} \Delta t \tag{11}$$

and analogously for RH. To distinguish "rare but severe" from "frequent but mild" violations, we report the average violation magnitude

$$\bar{\delta}^{(s)}_T = \frac{D^{(s)}_T}{\max\left(H^{(s)}_T, \varepsilon\right)}, \qquad \bar{\delta}^{(s)}_{RH} = \frac{D^{(s)}_{RH}}{\max\left(H^{(s)}_{RH}, \varepsilon\right)} \tag{12}$$

with a small $\varepsilon > 0$. CO_2 violations are computed as

$$D^+_{CO_2} = \sum_{k=1}^{N} \max\left(0, CO_{2,k} - CO_{2,\max}\right) \Delta t \tag{13}$$

$$H^+_{CO_2} = \sum_{k=1}^{N} \mathbb{I}\left\{CO_{2,k} > CO_{2,\max}\right\} \Delta t. \tag{14}$$

Using the above metrics, RH compliance improves in both severity and duration: D_{RH} decreases from 179.90 to 125.63% · h and H_{RH} decreases from 74.33 to 65.33 h (Table VII). Temperature compliance exhibits a tradeoff: the weighted severity decreases (D_T: 138.64 → 112.28 °C · h) and violations become milder on average ($\bar{\delta}_T$: 1.066 to 0.726 °C), while the violation duration increases (H_T: 130.00 to 154.67 h under the strict band $[T_{\min}, T_{\max}]$). This behavior is consistent with a "soft comfort" strategy that reduces excursion magnitude (degree-hours) without necessarily enforcing rapid re-entry into the strict band at every time step. Finally, CO_2 remains below $CO_{2,\max} = 1000$ ppm throughout (zero violation hours).

Change-effort (how much we perturb actuated variables): to quantify how strongly the algorithm modifies an actuated variable x relative to baseline, we report

$$J_{\Delta x} = \sum_{k=1}^{N} \left|x^{(1)}_k - x^{(0)}_k\right| \Delta t \tag{15}$$

which is an effort indicator (not a "better/worse" metric by itself). In our data, $J_{\Delta T} = 202.14\ ^\circ\mathrm{C}\cdot h$ and $J_{\Delta\mathrm{RH}} = 119.47\%\cdot h$, corresponding to average absolute adjustments of approximately 0.80 °C and 0.47% RH over the full horizon (253.67 h). The counterfactual improvements are obtained with moderate average adjustments (≈0.80 °C and ≈0.47% RH), indicating that the optimized trajectory does not require aggressive actuation.

VIII. IMPLEMENTATION

This section outlines how the proposed optimization can be implemented using an IoT-enabled actuation system. The main building blocks of the approach illustrated in Fig. 11. The approach can be implemented in two phases.

First, data are acquired from the indoor/outdoor environments for a suitable period of time (e.g., during a specific

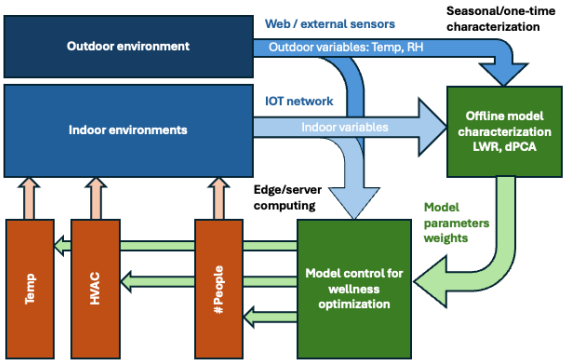

Fig. 11. Possible implementation of the optimization approach.

season period) by an IoT sensor network, in order to extract correlations. Autoscaling is performed, and an LWR model is constructed offline on a variable-dimensionality reduction obtained via dPCA on the time-series data. The elaboration identifies a number A of PCs, on which the dataset variable space is reduced, and the LWR kernel weights are constructed. The model classifies the wellness of each observation, either as the overall wellness average of the smart building or as a wellness metric of single environments in the case of distributed optimization. Then, a selection of observation targets $\mathbf{X}_{\text{ref}}$ with maximum indoor satisfaction is identified, e.g., using a threshold on the wellness metric. In the case of multiple environments, different thresholds and targets can be selected.

Second, the LWR kernel weights are loaded onto a server or a local device for edge-computing optimization in the runtime pipeline of the devices. Even if the process appears computationally intensive, the dimensionality reduction allows easy implementation on last-generation microcontrollers. In this stage, the type of actuation for feedback control must be defined. The optimization determines all variable changes allowed to vary within predefined boundaries, as indicated by (3). In particular, the runtime optimizer does not output an abstract score alone: it returns a corrected physical feature vector $\mathbf{x}_{\text{new}} = \mathbf{x} + \mathbf{\Delta}^{\star}$ that satisfies the imposed bounds. Actionable HVAC commands are then obtained by extracting the controllable components of $\mathbf{x}_{\text{new}}$ through the local BMS/controller interface (e.g., zone temperature setpoint and, when available, ventilation/humidity-related commands), while enforcing actuator saturation and optional rate limits. Nonactuated variables can be kept fixed by constraining their corresponding entries in $\mathbf{\Delta}^{\star}$ to zero.

The online block (LWR score computation followed by an SQP update) was profiled per sample. The measured latency is tightly concentrated around ~50–56 ms (median 53.14 ms and maximum 78.00 ms), with the runtime dominated by the LWR step and only a small additional overhead from SQP (median 3.07 ms). To provide a hardware-agnostic indicator, the real-time factor (RTF) is reported relative to the sampling interval $\Delta t = 20$ min: RTF $\approx 4.4 \times 10^{-5}$ at the median latency (and $\approx 6.5 \times 10^{-5}$ at worst case), indicating that the online computation is comfortably real time at the considered sampling rate. Compact memory footprints for the deployed parameters and an operation-count proxy are also reported as lightweight indicators for edge deployment. In the present campaign, the Tauanito INAIR device was used as the sensing and communication node. The runtime measurements reported above refer to the online dPCA–LWR–SQP computational pipeline profiled per sample, rather than to a closed-loop firmware implementation running directly on the INAIR MCU. Accordingly, these measurements demonstrate real-time computational feasibility at the adopted sampling rate, while direct MCU firmware deployment and closed-loop BMS validation are left for future work.

The runtime memory requirement of the online block can be expressed analytically and is therefore hardware-independent. In the reduced (score) space, storage is mainly due to the dPCA loadings, scaling parameters, a fixed-size lag buffer of $(l + 1)$ samples, and the LWR reference score set used for neighbor search. For the configuration used in this study ($K = 10$, $A = 7$, $\ell = 4$, and $N_{\text{ref}} = 344$), the total online memory is approximately 9.98 kB in single precision (or 19.96 kB in double precision), and it scales linearly with N_{ref}. Memory can be bounded by storing a compressed reference set (prototypes) instead of all reference points.

Depending on the actuation system, variables that are not feedback-actuated can be bounded, while actuated variables remain free. An example with Temp-only actuation was presented in Section VI-B.

IX. Conclusion

This work presents a framework for optimizing indoor environmental wellness by integrating IoT sensing, dynamic MVA, and nonlinear constrained optimization. By exploiting the temporal dependencies of environmental variables through dPCA, the proposed method achieves improved prediction accuracy compared to static approaches, particularly in nonlinear regression scenarios. The use of SQP within a reduced-rank score space allows the derivation of interpretable and feasible control actions while respecting physical and operational constraints.

Experimental validation on real-world data from three indoor environments demonstrated substantial improvements in predicted occupant satisfaction: from 5.18% to 94.81% in Room R_A, from 12.69% to 90.47% in R_B, and from 58.48% to 72.84% in R_C. The smaller gain in RC highlights the consistency of the method, as this room already exhibited high satisfaction levels before optimization and therefore required minimal environmental adjustment. These results confirm that the proposed approach can effectively identify and guide changes to environmental variables within operational limits, leading to significant comfort improvements.

Recent lightweight attention-based architectures provide a useful deployment-oriented analogy for our framework. For example, Usman et al. [64] proposed a lightweight transformer-driven multiscale trapezoidal attention network that emphasizes informative features through multiscale refinement and compact attention while controlling computational cost. Although this method is designed for image-based saliency detection and is not directly applicable to tabular HVAC/IoT time series, its principle is relevant to smart-building implementation: useful signals should be

emphasized without increasing the computational burden. In our framework, this role is achieved through dPCA loadings, reduced score-space modeling, and local LWR neighborhoods, which identify dominant variable relationships and influential operating regimes while preserving low memory/runtime requirements for edge and BMS deployment.

Future developments of this research will focus on two main directions. First, the implementation of the optimization algorithm on resource-constrained MCUs, enabling real-time, edge-computing deployment in smart building environments without relying on high-performance computing infrastructure. Second, the extension of the optimization problem to include energetic constraints, such as a maximum allowable energy consumption or operational budget, thus allowing the control strategy to explicitly trade off wellness improvements against energy savings and associated economic costs. These enhancements will further increase the applicability of the framework to practical, cost-aware building management scenarios.

X. Limitations and Future Work

This study demonstrates the proposed pipeline on a real-world dataset with occupant-reported TSVs; however, several limitations should be noted and motivate future extensions. The proposed optimization focuses on thermal sensation/satisfaction under feasibility constraints, while other wellness dimensions (e.g., IAQ, lighting, and acoustics) and energy are not jointly optimized due to limited metering and actuation interfaces in the current setting. Future work will extend the formulation to a multiobjective framework and evaluate scalability to multizone/building-level operation by accounting for shared HVAC constraints and interzone coupling. Because the present dataset contains aggregated room-level feedback, the model does not capture individual thermal preference variability. Personalized comfort models, occupant clustering, and sensitivity analysis under alternative comfort formulations, such as TSV-based, Fanger PMV/PPD-based, or hybrid satisfaction-based objectives, require additional occupant-level feedback and/or sensing of clothing and metabolic-rate proxies. Besides, the "PMV" variable used here is an occupant-reported TSV (ASHRAE 7-point scale), not a Fanger-computed PMV. Future extensions will compute Fanger PMV/PPD (with appropriate assumptions or additional sensing for clothing/metabolic proxies) and assess how optimization outcomes change under TSV- versus PMV-based comfort formulations.

In addition, our evaluation is based on a winter monitoring campaign and an offline counterfactual optimization, since closed-loop HVAC/BMS actuation was not available during data collection. Future work will validate robustness across multiple seasons and buildings and will address deployment aspects (sensor placement/reliability, edge/server execution, and BMS integration with actuator limits and logging) to enable closed-loop testing with actuator trajectories and, when available, metered energy.

As the dataset covers a single late-autumn/winter monitoring period, the present study cannot fully characterize seasonal effects or long-term changes in occupancy and weather conditions. Future work will extend the campaign to multiseason data collection and use strictly chronological validation protocols to assess robustness under different outdoor conditions, occupancy schedules, and building-operation regimes. While the dataset covers three rooms with clearly different occupancy regimes, it remains limited in duration and spatial diversity. An extended data collection plan is being considered to include longer monitoring periods and additional buildings/seasons, in order to strengthen external validity and evaluate cross-building generalization.

Appendix A
Fanger's Model

The concept of thermal comfort is influenced by multiple environmental, physiological, and psychological factors. It plays a critical role in occupant well-being and productivity, particularly in indoor environments, such as homes, offices, and public buildings. To standardize comfort assessment, international regulations have been established, including ISO 7730 [65] for Europe and ASHRAE 55 [66] for the United States. Both rely on Fanger's model [17], developed in 1970, which evaluates comfort under steady-state and uniform environmental conditions based on physical variables, such as air Temp, RH, air velocity, metabolic rate, and clothing insulation. A detailed application of this model in an IoT context can be found in [26] and [67].

Fanger's model characterizes the distribution of indoor comfort for a population by relating the PMV index to the predicted percentage of dissatisfied people (PPD). The PMV represents the average thermal sensation of a large group of individuals on a seven-point scale ranging from +3 (hot) to −3 (cold), with 0 indicating a neutral sensation. The PPD curve reaches its minimum (about 5% dissatisfied) at PMV = 0 and increases symmetrically as PMV moves toward hot or cold extremes. As shown in Fig. 10, the percentage of dissatisfied people approaches 100% asymptotically as temperature conditions reach extreme values.

Appendix B
Optimization by SQP

A key point of the optimization algorithm is to work in the reduced-rank score space, which greatly reduces the number of variables and matrix sizes. Therefore, any new or dataset observations should be projected into the score space. Referring to paper's notation, the score vector $\mathbf{T}'_{\text{obj}} \in \mathbb{R}^{1\times A}$ is linearly updated when its corresponding observation $\mathbf{X}_{\text{obj}} \in \mathbb{R}^{1\times K}$ changes by a vector $\boldsymbol{\Delta} \in \mathbb{R}^{1\times K}$. The objective is to compute the optimal $\boldsymbol{\Delta}$ that shifts the score $\mathbf{T}'_{\text{obj}} \in \mathbb{R}^{1\times A}$ as close as possible to the mean reference score $\mathbf{T}_{\text{ref}} \in \mathbb{R}^{1\times A}$. Using the equivalence

$$\begin{aligned}\mathbf{T}'_{\text{obj}}(\boldsymbol{\Delta}) &= (\mathbf{X}_{\text{obj}} + \boldsymbol{\Delta}) \cdot \mathbf{L} \\ &= \mathbf{X}_{\text{obj}} \cdot \mathbf{L} + \boldsymbol{\Delta} \cdot \mathbf{L} \\ &= \mathbf{T}_{\text{obj}} + \boldsymbol{\Delta} \cdot \mathbf{L} \end{aligned} \tag{16}$$

where $\mathbf{L}$ is the loading matrix obtained from PCA, $\mathbf{T}_{\text{obj}}$ is the original score, and we seek the variation Δ that minimizes the Euclidean distance in score space to T_{ref} while penalizing

large deviations in the original variable space. This leads to the regularized problem

$$\arg\min_{\boldsymbol{\Delta}} \ \left\|\boldsymbol{T}_{\text{ref}} - \left(\boldsymbol{T}_{\text{obj}} + \boldsymbol{\Delta}L\right)\right\|_2^2 + \lambda\left\|\boldsymbol{\Delta}\right\|_2^2 \tag{17}$$

subject to the physical constraints of (3). Here, $\lambda > 0$ is the regularization parameter; we set $\lambda = 1$ in our case. If the constraints are ignored, (17) reduces to a standard Tikhonov-regularized least-squares problem

$$\arg\min_{\boldsymbol{x}} \ \left\|\boldsymbol{A}\boldsymbol{x} - \boldsymbol{b}\right\|_2^2 + \lambda\left\|\boldsymbol{x}\right\|_2^2 \tag{18}$$

where

$$\begin{aligned}\boldsymbol{A} &= L^\top \in \mathbb{R}^{A\times K}\\ \boldsymbol{b} &= \left(\boldsymbol{T}_{\text{ref}} - \boldsymbol{T}_{\text{obj}}\right)^\top \in \mathbb{R}^{A\times 1}\\ \boldsymbol{x} &= \boldsymbol{\Delta}^\top \in \mathbb{R}^{K\times 1}.\end{aligned}$$

However, when the physical bounds above are enforced, the problem becomes a nonlinear constrained optimization, becoming part of a nonlinear programming (NLP) problem.

NLP is a classical approach for finding local minima or maxima of a scalar-valued objective function $f(\boldsymbol{x})$ defined on $\mathbb{R}^K$, subject to equality and inequality constraints. The optimization problem can be formalized as

$$\begin{aligned}\boldsymbol{x} = \ &\arg\min_{\boldsymbol{x}} f(\boldsymbol{x})\\ &\text{s.t.} \quad g(\boldsymbol{x}) = 0\\ &\qquad h(\boldsymbol{x}) \geq 0 \rightarrow h(\boldsymbol{x}) - s^2 = 0\\ &\qquad k(\boldsymbol{x}) \leq 0 \rightarrow k(\boldsymbol{x}) + t^2 = 0.\end{aligned}$$

Here, s and t are the slack variables introduced to handle inequality constraints. The constrained problem is transformed into an unconstrained one by defining the Lagrangian function $\mathcal{L}(\boldsymbol{x}, \lambda, \theta, \phi)$ as

$$\begin{aligned}&\mathcal{L}(\boldsymbol{x}, \lambda, \theta, \phi)\\ &= f(\boldsymbol{x}) + \lambda g(\boldsymbol{x}) - \theta\left[h(\boldsymbol{x}) - s^2\right] + \phi\left[k(\boldsymbol{x}) + t^2\right]\end{aligned} \tag{19}$$

where λ, θ, and ϕ are the Lagrange multipliers associated with the constraints.

To satisfy the necessary conditions for optimality, the gradient of the Lagrangian must be zero, leading to the system

$$\begin{aligned}&\nabla\mathcal{L} = 0\\ &\Rightarrow \text{w.r.t. } \boldsymbol{x} : \nabla f(\boldsymbol{x}) + \lambda\nabla g(\boldsymbol{x}) - \theta\nabla h(\boldsymbol{x}) + \phi\nabla k(\boldsymbol{x}) = 0\\ &\text{w.r.t. } \lambda : g(\boldsymbol{x}) = 0\\ &\text{w.r.t. } \theta : h(\boldsymbol{x}) = s^2\\ &\text{w.r.t. } \phi : k(\boldsymbol{x}) = -t^2.\end{aligned} \tag{20}$$

This implies that the gradient of the objective function $f(\boldsymbol{x})$ is aligned with the gradients of the constraint functions. The optimization is carried out by solving this system, typically using numerical methods, such as SQP. In general, the function $f(\boldsymbol{\Delta})$ is minimized under constraints $g(\boldsymbol{\Delta}) = 0$. Direct Newton's method for finding critical points employs inverse of Hessian matrix of the Lagrangian $\nabla^2\mathcal{L}$. However, it is typically ill-conditioned or near-singular, leading the convergence procedure critical. SQP overcomes these problems by finding a search direction $\boldsymbol{\delta}_k$ that minimizes the second-order Taylor expansion

$$\begin{aligned}&\arg\min_{\boldsymbol{\delta}_k} \quad f(\boldsymbol{\Delta}_k) + \nabla f(\boldsymbol{\Delta}_k)'\,\boldsymbol{\delta}_k + \frac{1}{2}\boldsymbol{\delta}_k'\nabla^2\mathcal{L}(\boldsymbol{\Delta}_k, \lambda)\,\boldsymbol{\delta}_k\\ &\text{s.t.} \quad g(\boldsymbol{\Delta}_k) + \nabla g(\boldsymbol{\Delta}_k)'\,\boldsymbol{\delta}_k = 0\end{aligned} \tag{21}$$

that does not need Hessian inversion. The SQP algorithm begins with an initial point $(\boldsymbol{\Delta}_0, \lambda_0)$ and iteratively updates the solution using

$$\left[\boldsymbol{\Delta}_{k+1}, \lambda_{k+1}\right] = \left[\boldsymbol{\Delta}_k, \lambda_k\right] + \boldsymbol{\delta}_k. \tag{22}$$

This process continues until the convergence criteria are met.

ACKNOWLEDGMENT

Views and opinions expressed are, however, those of the authors only and do not necessarily reflect those of the European Union or Chips JU.

REFERENCES

[1] J. Ngarambe, G. Y. Yun, and M. Santamouris, "The use of artificial intelligence (AI) methods in the prediction of thermal comfort in buildings: Energy implications of AI-based thermal comfort controls," *Energy Buildings*, vol. 211, Mar. 2020, Art. no. 109807.

[2] L. A. López-Pérez and J. J. Flores-Prieto, "Adaptive thermal comfort approach to save energy in tropical climate educational building by artificial intelligence," *Energy*, vol. 263, Jan. 2023, Art. no. 125706.

[3] N. L. Rane, S. P. Choudhary, and J. Rane, "Enhancing thermal comfort through leading-edge design, monitoring, and optimization technologies: A review," *Sustain. Clean Buildings*, vol. 1, pp. 123–146, Sep. 2024.

[4] W. Hu, Y. Wen, K. Guan, G. Jin, and K. J. Tseng, "ITCM: Toward learning-based thermal comfort modeling via pervasive sensing for smart buildings," *IEEE Internet Things J.*, vol. 5, no. 5, pp. 4164–4177, Oct. 2018.

[5] Z. Zhang and K. P. Lam, "Practical implementation and evaluation of deep reinforcement learning control for a radiant heating system," in *Proc. 5th Conf. Syst. Built Environments*, Shenzen, China, Nov. 2018, pp. 148–157.

[6] F. Cicirelli, A. Guerrieri, C. Mastroianni, G. Spezzano, and A. Vinci, "Thermal comfort management leveraging deep reinforcement learning and human-in-the-loop," in *Proc. IEEE Int. Conf. Hum.-Mach. Syst. (ICHMS)*, Sep. 2020, pp. 1–6.

[7] G. Gao, J. Li, and Y. Wen, "Deepcomfort: Energy-efficient thermal comfort control in buildings via reinforcement learning," *IEEE Internet Things J.*, vol. 7, no. 9, pp. 8472–8484, Sep. 2020.

[8] M. Al-Adhami, S. Wu, and E. Delzendeh, "A real-time approach to evaluate occupants'thermal comfort in the indoor environment," in *Proc. 37th CIB W78 Inf. Technol. Construction Conf.* (CIB W), 2020, pp. 477–487.

[9] F. Alsaleem, M. K. Tesfay, M. Rafaie, K. Sinkar, D. Besarla, and P. Arunasalam, "An IoT framework for modeling and controlling thermal comfort in buildings," *Frontiers Built Environ.*, vol. 6, p. 87, Jun. 2020.

[10] L. Yu et al., "Multi-agent deep reinforcement learning for HVAC control in commercial buildings," *IEEE Trans. Smart Grid*, vol. 12, no. 1, pp. 407–419, Jan. 2021.

[11] A. Naeem et al., "Energyplus as a computational engine for commercial building operational digital twins," *Energy Buildings*, vol. 329, Feb. 2025, Art. no. 115257.

[12] W. Zhang, Y. Wen, K. J. Tseng, and G. Jin, "Demystifying thermal comfort in smart buildings: An interpretable machine learning approach," *IEEE Internet Things J.*, vol. 8, no. 10, pp. 8021–8031, May 2020.

[13] B. Chen, Z. Cai, and M. Bergés, "Gnu-RL: A precocial reinforcement learning solution for building HVAC control using a differentiable MPC policy," in *Proc. 6th ACM Int. Conf. Syst. Energy-Efficient Buildings*, Cities, Transp., Nov. 2019, pp. 316–325.

[14] K. Kurte et al., "Evaluating the adaptability of reinforcement learning based HVAC control for residential houses," *Sustainability*, vol. 12, no. 18, p. 7727, Sep. 2020.

[15] W. Jiang, D. Meng, and W. Jin, "Comfort management based on predictive PMV in smart homes," in *Proc. 4th Int. Conf. Electron. Inf. Eng. Comput. Sci. (EIECS)*. Yanji, China: IEEE, Sep. 2024, pp. 929–932.

[16] W. Xue, H. Wang, and K. Li, "PMV inverse model-based indoor thermal environment control for thermal comfort and energy saving," in *Proc. 41st Chin. Control Conf. (CCC)*. Hefei, China: IEEE, Jul. 2022, pp. 5294–5299.

[17] P. O. Fanger, *Thermal Comfort. Analysis and Applications in Environmental Engineering*. Copenhagen, Denmark: Danish Technical Press, 1970.

[18] A. Berardinelli, G. Luciani, M. Crescentini, A. Romani, M. Tartagni, and L. Ragni, "Application of non-linear statistical tools to a novel microwave dipole antenna moisture soil sensor," *Sensors Actuators A: Phys.*, vol. 282, pp. 1–8, Oct. 2018.

[19] L. Franceschelli et al., "Non-intrusive microwave technique for direct detection of concrete compressive strength monitoring by multivariate modeling," *Meas., J. Int. Meas. Confederation*, vol. 206, pp. 1–10, Jan. 2023.

[20] O. Afif et al., "A versatile, machine-learning-enhanced RF spectral sensor for developing a trunk hydration monitoring system in smart agriculture," *Sensors*, vol. 24, no. 19, p. 6199, Sep. 2024.

[21] S. Wang and F. Xiao, "AHU sensor fault diagnosis using principal component analysis method," *Energy Buildings*, vol. 36, no. 2, pp. 147–160, Feb. 2004.

[22] L. Burgas, J. Colomer, J. Meléndez, F. I. Gamero, and S. Herraiz, "Integrated unfold-PCA monitoring application for smart buildings: An AHU application example," *Energies*, vol. 14, no. 1, p. 235, 2021.

[23] C. Manna, N. Wilson, and K. N. Brown, "Personalized thermal comfort forecasting for smart buildings via locally weighted regression with adaptive bandwidth," in *Proc. 2nd Int. Conf. Smart Grids Green IT Syst.*, 2013, pp. 32–40.

[24] D. Schwingshackl, J. Rehrl, and M. Horn, "LoLiMoT based MPC for air handling units in HVAC systems," *Building Environ.*, vol. 96, pp. 250–259, Feb. 2016.

[25] E. T. Maddalena, Y. Lian, and C. N. Jones, "Data-driven methods for building control—A review and promising future directions," *Control Eng. Pract.*, vol. 95, Feb. 2020, Art. no. 104211.

[26] O. Afif, G. Ingenito, M. Bellomo, A. Romagnoli, M. Tartagni, and A. Romani, "IoT-based indoor thermal comfort prediction using multivariate statistical analysis," *IEEE Sensors J.*, vol. 25, no. 8, pp. 14358–14369, Apr. 2025.

[27] (2024). *Visual Crossing Weather Data Services*. [Online]. Available: https://www.visualcrossing.com/

[28] L. Gugliermetti, F. Cumo, and S. Agostinelli, "A future direction of machine learning for building energy management: Interpretable models," *Energies*, vol. 17, no. 3, p. 700, Feb. 2024.

[29] Z. Chen, F. Xiao, F. Guo, and J. Yan, "Interpretable machine learning for building energy management: A state-of-the-art review," *Adv. Appl. Energy*, vol. 9, Feb. 2023, Art. no. 100123.

[30] Z. An, X. Ding, and W. Du, "Go beyond black-box policies: Rethinking the design of learning agent for interpretable and verifiable hvac control," in *Proc. 61st ACM/IEEE Design Autom. Conf.*, Nov. 2024, pp. 1–6.

[31] S. Cho and C. S. Park, "Rule reduction for control of a building cooling system using explainable AI," *J. Building Perform. Simul.*, vol. 15, no. 6, pp. 832–847, 2022.

[32] A. Stieger, C. Stippel, A. Bratukhin, R. Hoch, and T. Sauter, "How many sensors are necessary?—Sensor reduction in HVAC control modeling using feature attribution and rule extraction," *IEEE Sensors Lett.*, vol. 9, no. 10, pp. 1–4, Oct. 2025.

[33] J. Goldfeder and J. Sipple, "Reducing carbon emissions at scale: Interpretable and efficient to implement reinforcement learning via policy extraction," in *Proc. 11th ACM Int. Conf. Syst. Energy-Efficient Buildings*, Cities, Transp., Oct. 2024, pp. 403–407.

[34] O. Kotevska et al., "Methodology for interpretable reinforcement learning model for HVAC energy control," in *Proc. IEEE Int. Conf. Big Data (Big Data)*, Atlanta, GA, USA, Dec. 2020, pp. 1555–1564.

[35] L. Zhang and Z. Chen, "Large language model-based interpretable machine learning control in building energy systems," *Energy Buildings*, vol. 313, Jun. 2024, Art. no. 114278.

[36] Y. Dong and S. J. Qin, "A novel dynamic pca algorithm for dynamic data modeling and process monitoring," *J. Process Control*, vol. 67, pp. 1–11, Jul. 2018.

[37] C. G. Atkeson, A. W. Moore, and S. Schaal, "Locally weighted learning," *Artif. Intell. Rev.*, vol. 11, pp. 11–73, Feb. 1997.

[38] D. Perez-Guaita et al., "Modified locally weighted—Partial least squares regression improving clinical predictions from infrared spectra of human serum samples," *Talanta*, vol. 107, pp. 368–375, Mar. 2013.

[39] E. E. Broday, J. A. Moreto, A. A. de Paula Xavier, and R. de Oliveira, "The approximation between thermal sensation votes (TSV) and predicted mean vote (PMV): A comparative analysis," *Int. J. Ind. Ergonom.*, vol. 69, pp. 1–8, Jan. 2019.

[40] R. Bro and A. K. Smilde, "Principal component analysis," *Anal. Methods*, vol. 6, no. 9, pp. 2812–2831, 2014.

[41] S. Wold, "Cross-validatory estimation of the number of components in factor and principal components models," *Technometrics*, vol. 20, no. 4, pp. 397–405, Nov. 1978.

[42] W. Ku, R. H. Storer, and C. Georgakis, "Disturbance detection and isolation by dynamic principal component analysis," *Chemometric Intell. Lab. Syst.*, vol. 30, no. 1, pp. 179–196, Nov. 1995.

[43] T. Naes, T. Isaksson, and B. Kowalski, "Locally weighted regression and scatter correction for near-infrared reflectance data," *Anal. Chem.*, vol. 62, no. 7, pp. 664–673, Apr. 1990.

[44] V. Centner and D. L. Massart, "Optimization in locally weighted regression," *Anal. Chem.*, vol. 70, no. 19, pp. 4206–4211, Oct. 1998.

[45] S. S. Saha, S. S. Sandha, and M. Srivastava, "Machine learning for microcontroller-class hardware: A review," *IEEE Sensors J.*, vol. 22, no. 22, pp. 21362–21390, Nov. 2022.

[46] S. Branco, A. G. Ferreira, and J. Cabral, "Machine learning in resource-scarce embedded systems, FPGAs, and end-devices: A survey," *Electronics*, vol. 8, no. 11, p. 1289, Nov. 2019.

[47] J.-A. Ting, F. Meier, S. Vijayakumar, and S. Schaal, "*Locally weighted regression for control*," in *Encyclopedia of Machine Learning and Data Mining*. Cham, Switzerland: Springer, 2016, pp. 1–14.

[48] D. R. Brillinger, *Time Series: Data Analysis and Theory*. Philadelphia, PA, USA: SIAM, 2001.

[49] C. Liu, J. Bai, and F. Wu, "Fault diagnosis using dynamic principal component analysis and GA feature selection modeling for industrial processes," *Processes*, vol. 10, no. 12, p. 2570, Dec. 2022.

[50] H. Khan, M. T. Usman, I. Rida, and J. Koo, "Attention enhanced machine instinctive vision with human-inspired saliency detection," *Image Vis. Comput.*, vol. 152, Dec. 2024, Art. no. 105308.

[51] J. H. Stock and M. W. Watson, "Forecasting using principal components from a large number of predictors," *J. Amer. Stat. Assoc.*, vol. 97, no. 460, pp. 1167–1179, 2002.

[52] C. Schumacher, "Forecasting German GDP using alternative factor models based on large datasets," *J. Forecasting*, vol. 26, no. 4, pp. 271–302, Jul. 2007.

[53] S. Siami-Namini, N. Tavakoli, and A. S. Namin, "A comparative analysis of forecasting financial time series using arima, LSTM, and BILSTM," 2019, *arXiv:1911.09512*.

[54] S. Siami-Namini, N. Tavakoli, and A. Siami Namin, "A comparison of ARIMA and LSTM in forecasting time series," in *Proc. 17th IEEE Int. Conf. Mach. Learn. Appl. (ICMLA)*, Dec. 2018, pp. 1394–1401.

[55] J. Durbin and S. J. Koopman, *Time Series Analysis By State Space Methods*. London, U.K.: Oxford Univ. Press, 2012.

[56] P. T. Boggs and J. W. Tolle, "Sequential quadratic programming," *Acta Numerica*, vol. 4, pp. 1–51, Jan. 1995.

[57] J. Nocedal and S. J. Wright, *Numerical Optimization*, 2nd ed., New York, NY, USA: Springer, 2006.

[58] F. Wang, S. Chawla, and W. Liu, "Tikhonov or lasso regularization: Which is better and when," in *Proc. IEEE 25th Int. Conf. Tools Artif. Intell.*, Nov. 2013, pp. 795–802.

[59] A. Franco, C. Bartoli, P. Conti, L. Miserocchi, and D. Testi, "Multi-objective optimization of HVAC operation for balancing energy use and occupant comfort in educational buildings," *Energies*, vol. 14, no. 10, p. 2847, May 2021.

[60] A. Aswani, N. Master, J. Taneja, A. Krioukov, D. Culler, and C. Tomlin, "Quantitative methods for comparing different HVAC control schemes," in *Proc. 6th Int. ICST Conf. Perform. Eval. Methodologies Tools*, Oct. 2012, pp. 326–332.

[61] W. T. Grondzik, *Air-conditioning System Design Manual*. Amsterdam, The Netherlands: Elsevier, 2007.

[62] R. Zhu, J. Yang, J. Yang, and H. Yuan, "Study on the influence of interior design temperature on energy consumption and carbon emission of large space buildings in cold regions," in *Proc. 7th Int. Conf. Smart Grid Smart Cities (ICSGSC)*, Lanzhou, China, Sep. 2023, pp. 624–628.

[63] P. Carrer et al., "On the development of health-based ventilation guidelines: principles and framework," *Int. J. Environ. Res. Public Health*, vol. 15, no. 7, p. 1360, 2018.

[64] M. T. Usman, H. Khan, I. Rida, and J. Koo, "Lightweight transformer-driven multi-scale trapezoidal attention network for saliency detection," *Eng. Appl. Artif. Intell.*, vol. 155, Sep. 2025, Art. no. 110917.

[65] B. W. Olesen and K. C. Parsons, "Introduction to thermal comfort standards and to the proposed new version of EN ISO 7730," *Energy Buildings*, vol. 34, no. 6, pp. 537–548, Jul. 2002.
[66] B. W. Olesen and G. S. Brager, "A better way to predict comfort: The new ASHRAE standard 55–2004," Tech. Rep., Aug. 2004.
[67] O. Afif, G. Ingenito, M. Bellomo, M. Tartagni, and A. Romani, "IoT-driven indoor wellness optimization using a multivariate analysis algorithm," in *Proc. IEEE 11th World Forum Internet Things (WF-IoT)*, Oct. 2025, pp. 1–6.

Gaetano Ingenito received the master's degree in biomedical engineering from the Alma Mater Studiorum–Universitá di Bologna, Cesena, Italy, in March 2024.

He is currently working as a Data Analyst with Vantive Corporate, Medolla, Italy, in Kidney Care, covering the roles of Eng, QS, and Data Trending. His work comprehends the use of statistical and programming tools for complaint analysis in quality decisions. His master's thesis, titled "Indoor Wellness: A Machine Learning Technique for Internet of Things Sensor Data," was conducted at the Alma Mater Studiorum–Universitá di Bologna.

Oumaima Afif (Graduate Student Member, IEEE) received the master's degree in aerospace engineering from the School of Aerospace, Université Internationale de Rabat (UIR), Rabat, Morocco, in July 2022. She is currently pursuing the Ph.D. degree with the Department of Electrical, Electronic, and Information Engineering "Guglielmo Marconi" (DEI), University of Bologna, Cesena, Italy.

Her master's thesis, titled "Routing in CubeSat Based on Architecture and Topology Design," was conducted at Ca' Foscari University of Venice, Venice, Italy. Her research interests include the application of machine learning (ML) in autonomous sensing systems, with a focus on analyzing ML predictive models based on statistical paradigms (e.g., multivariate analysis methods).

Aldo Romani (Member, IEEE) received the Dr.Eng. degree in electronic engineering and the Ph.D. degree in electronics engineering, computer science, and telecommunications from the University of Bologna, Bologna, Italy, in 2001 and 2005, respectively.

In 2008, he joined the University of Bologna, Cesena, Italy, where he is currently an Associate Professor. He has been involved with various applications of sensors and circuits for energy harvesting. He has authored or co-authored more than 120 international scientific publications.

Dr. Romani was a co-recipient of the Jan Van Vessem Award at the 2004 IEEE International Solid-State Circuits Conference.

Marco Tartagni (Member, IEEE) received the M.S. degree in EE and the Ph.D. degree in EECS from the University of Bologna, Bologna, Italy, in 1988 and 1992, respectively.

In 1992, he joined the Department of EE, Caltech, Pasadena, CA, USA, initially as a Visiting Student and later in 1994 as a Research Fellow. During his time there, he focused on various aspects of analog very large-scale integration (VLSI) for image processing. Since 1995, he has been a Faculty Member with the Department of EE, University of Bologna, where he currently holds the position of a Full Professor. From 1996 to 2001, he was the Team Leader with the Collaborative Laboratory, Bologna, between STMicroelectronics and the University of Bologna. His work during this period involved the development of intelligent CMOS sensors, including CMOS cameras and the world's first prototype of a silicon-based fingerprint capacitive sensor. He is currently focusing his research on sensor design theory and machine learning-assisted spectral sensing.

Dr. Tartagni was a co-recipient of IEEE Van Vessem Outstanding Paper Award, which he received at the 2004 ISSCC Conference. He has also taken on the role of European Coordinator for the FP6 Receptronics project in the field of nanoelectronics. Furthermore, he has served on the Scientific Committees for the IEEE Custom Integrated Circuit Conference (CICC) from 2017 to 2021 and the IEEE International Electron Device Meeting (IEDM) in 2023.