Title

A Probabilistic Approach to Load Modeling for Central HVAC Systems in Large Commercial Buildings for Retrofit Decisions Under Uncertainty

Journal

ASHRAE Transactions, 132(1)

ISSN

0001-2505

Authors

Ham, SW

Excell, L

Kim, D

Publication Date

2026-04-06

DOI

10.63044/w26ham179

A Probabilistic Approach to Load Modeling for Central HVAC Systems in Large Commercial Buildings for Retrofit Decisions Under Uncertainty

Sang woo Ham, PhD
Member ASHRAE

Lauren Excell, PhD
Affiliate Member ASHRAE

Donghun Kim, PhD
Member ASHRAE

ABSTRACT

Retrofitting central HVAC systems in large commercial buildings with advanced technologies like heat recovery chillers (HRCs) offers a significant opportunity to enhance energy efficiency. However, analyzing these retrofits is challenging with traditional whole-building simulation tools, which require intensive calibration and struggle to model innovative system configurations and controls. To overcome these limitations, this study proposes a load profile-based retrofit analysis framework that provides better decisions under uncertainty. The main focus of this paper is the development of a probabilistic load profile model that can be used in the framework by using exploratory data analysis (EDA) of measured building data to properly quantify its inherent variability. A non-parametric Gaussian Process (GP) model was employed to capture the time- and weather-dependent characteristics of the heating load while explicitly modeling its uncertainty. The model's effectiveness is demonstrated through strong predictive performance on unseen data and physically interpretable insights into load behavior. This data-driven, probabilistic load profile serves as a robust and flexible input for subsequent system simulations, enabling a more confident and statistically sound analysis of retrofit potential.

INTRODUCTION

Heat pump technologies, specifically water-to-water heat recovery chillers (HRCs) (Carrier 2022), can produce hot water at temperatures suitable for building heating. This capability allows them to capture and repurpose waste heat from a facility's cooling process. HRCs work by simultaneously generating chilled water while transferring the heat removed from the cooling loop—plus the compressor work—to a separate hot water system. By repurposing this thermal energy, these systems can achieve a total coefficient of performance (COP) as high as 7.0, making them a highly efficient alternative to operating separate chillers and boilers (California Energy Design Assistance 2023).

Large commercial and institutional buildings, such as hospitals, hotels, and pharmaceutical manufacturing facilities, are ideal candidates for HRC systems (California Energy Design Assistance 2023). These facilities often require year-round, simultaneous heating, process water, and cooling—for example, to maintain precise humidity and temperature setpoints in GMP (Good Manufacturing Practices) spaces (Heemer, Mitrovic, and Scheer 2011) while also conditioning occupied areas. In such scenarios, the constant cooling load provides a consistent source of waste heat that can be recovered to meet ongoing heating demands for processes, space conditioning, or service water heating.

Retrofitting older, less efficient heating plants, such as those with conventional boilers, with HRC systems presents

Sang woo Ham is a Technology Researcher, **Lauren Excell** is a Postdoctoral researcher, and **Donghun Kim** is a Computational Research Scientist at Lawrence Berkeley National Lab, Berkeley, California.

a significant opportunity to improve overall plant efficiency and reduce primary energy consumption. However, the design and implementation of these advanced systems are far from straightforward. Engineers face numerous decisions regarding plant configurations—such as parallel or series arrangements—and the potential integration of thermal energy storage (TES) to decouple heating and cooling demands, which can lead to highly complex system designs and control sequences aimed at maximizing efficiency (Brandon Gill 2021).

The analysis of HVAC retrofits traditionally relies on year-round, whole-building energy simulations (Lee et al. 2015). These simulations model the building's envelope, internal loads, and the specific HVAC system to generate a deterministic load profile, which then serves as the basis for evaluating the performance of proposed upgrades (Wilson et al. 2022). However, because whole-building energy simulations are quasi–steady-state in nature, they are limited to the simplified control strategies and plant configurations supported by the simulation engines.

Various approaches exist for conducting a plant-level analysis of HVAC and control retrofits. Some methods (Catrini et al. 2020) rely on integrating third-party simulation tools, which use building and historical data to evaluate different retrofit scenarios. Other approaches (Liang et al. 2025) use detailed mathematical models of plant components, calibrated with operational data, to predict the performance improvements from retrofitting. Regardless of the specific method, a focused, plant-level analysis is often required to accurately evaluate the performance of various HVAC retrofit scenarios. While fully coupled simulations that model both the building and its mechanical plant in detail are possible (Bhattacharya et al. 2021; Tommasi et al. 2018), they can introduce computational complexity, especially when comparing multiple retrofit options. To alleviate this complexity, a load-based approach decouples the building's thermal load from the performance of the plant (Hinkelman et al. 2022). By using a building load profile as the input to the plant model, this method allows for a detailed evaluation of various plant configurations and control strategies without the need to re-run the entire building simulation for each scenario. Dynamic simulation tools like Modelica (Wetter et al. 2014) are particularly well-suited for this type of plant-level analysis, providing the flexibility to model novel system configurations and to optimize sophisticated control sequences (Huang, Zuo, and Sohn 2016; Bai et al. 2024).

While this deterministic approach provides a valuable baseline, it does not fully capture the inherent variability and uncertainty of real-world building operations. Factors such as occupant behavior are highly stochastic, and modeling this uncertainty can provide deeper insights into a building's range of potential energy consumption patterns, leading to more robust and realistic predictions (Chen et al. 2022). Parametric analysis of whole-building simulation with stochastic simulation inputs can be used to capture a range of operating conditions, however the computational intensity of this approach limits it to the analysis of a handful of variables, limiting the range of operating scenarios represented. A data-driven approach can better capture the extrema of the uncertainty without requiring as high of computational intensity.

Just as uncertainty analysis has been incorporated into whole-building simulations to provide more reliable retrofit guidance, a similar probabilistic approach is needed for plant-level load profiles. The performance of different plant configurations can be highly sensitive to variations in the building load. By developing a probabilistic load profile that captures this uncertainty, we can quantify how different retrofit options will perform across a range of likely conditions, rather than relying on a single, deterministic outcome. This is a critical factor for decision-making, as it provides a clearer picture of the potential risks and rewards associated with a given retrofit investment. While uncertainties in plant-level measurements and equipment performance are an important consideration (Wang et al. 2023), the uncertainty originating from the building's load profile is often the major differentiating factor in a comparative retrofit analysis. This allows for a robust comparison of how different potential scenarios will perform under a realistic range of operating demands.

In this paper, we present a probabilistic load profile model specifically developed for analyzing HRC retrofit applications under uncertainty. Since the performance and economic viability of HRC systems are highly sensitive to building load variations, this probabilistic approach is essential for evaluating how different retrofit options will perform across a realistic range of conditions. This provides a more robust basis for investment decisions than an analysis based on a single, deterministic load profile. This paper begins by introducing the overall retrofit analysis framework, which uses the developed load profile as a primary input for a detailed HRC plant model developed in the dynamic simulation language Modelica. We then present the real-world heating load data and the Exploratory Data Analysis (EDA) used to inform the model's formulation by identifying key time- and weather-dependent characteristics. This analysis is crucial for separating deterministic, predictable patterns from true stochastic noise, which allows for the proper characterization of the model's uncertainty. Following this, we describe the development of the non-parametric GP model and how the model can capture the temporal and weather-specific information by analyzing kernels for the model validation. The paper concludes with a summary of the work and directions for future research.

FRAMEWORK FOR EVALUATING CENTRAL PLANT RETROFIT

Figure 1 presents a load-based dynamic simulation framework designed to evaluate central plant retrofit scenarios. This approach uses an annual building load profile, provided as a CSV input, to serve as a dynamic boundary condition for testing multiple potential retrofit configurations. These configurations, which include detailed models of equipment and control systems, are developed in Modelica to capture their complex physical behavior. By simulating the performance of each configuration against the measured load profile, key metrics such as annual energy consumption can be directly compared to identify the optimal design. This method avoids the need to create a detailed building model for each scenario, allowing for a focused and efficient analysis of plant-side retrofit options.

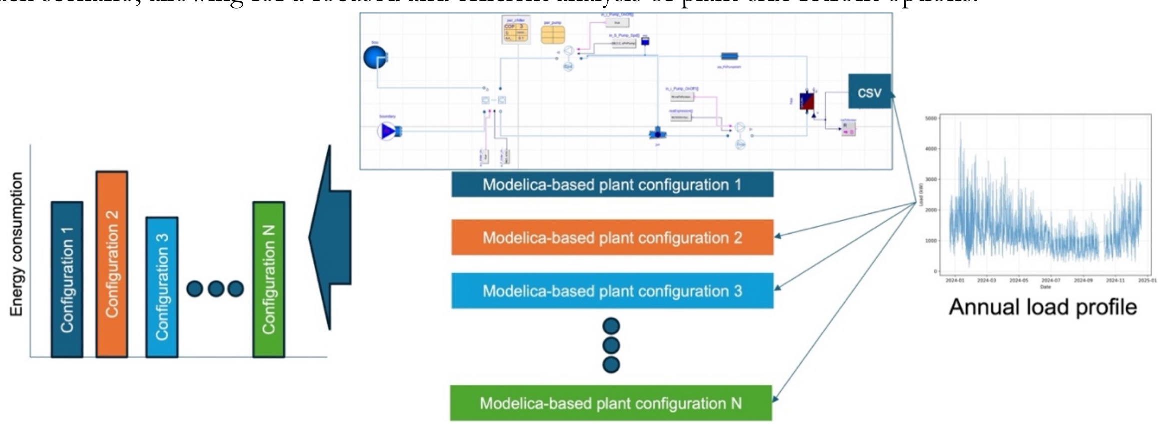

Figure 1 Overview of the load-based simulation framework for central plant retrofit analysis.

While the retrofit framework provides the overall context, the primary contribution of this paper is the development of the stochastic load profile that serves as its key input. Moving beyond a deterministic approach, a model that generates a stochastic load profile allows for more robust decision-making under real-world uncertainty. Therefore, we focus on creating a model that represents the inherent variability of a building's energy demands. By using this probabilistic load profile, the analysis framework can evaluate retrofit options against a realistic range of operating conditions, leading to more confident and reliable investment decisions.

LOAD-PROFILE MODEL

Building description

To develop a realistic load profile model, this study investigates the measured heating load profile from a representative facility to capture its operational dynamics. The building is a modern, mid-sized, mixed-use laboratory and office facility constructed in the mid-2000s, with a total floor area of approximately 150,000 ft^2 (13,935 m^2). The facility's space is highly varied, with laboratory areas comprising approximately 25% of the total floor area. Office and conference spaces account for another 16%, while a significant portion is dedicated to a dining hall and kitchen (13%). The remaining space is used for corridors, mechanical rooms, and other support areas. The building's heating is provided by a central plant with condensing boilers that serve hot water coils in both the main air handling units and the variable air volume (VAV) terminal boxes for zone-level reheat. These condensing boilers represent the primary target for a potential retrofit with a HRC system.

Exploratory data analysis (EDA)

To develop an accurate model, we first conducted an exploratory data analysis (EDA) on the measured year-round heating load profile to understand its underlying dynamics. The goal of our modeling approach is to capture the deterministic patterns based on known features (such as time and weather) as accurately as possible, so that the remaining variability can be modeled as the stochastic component. Decomposing a time-series dataset into trends at different timescales—such as weekly and seasonal—is a common and effective analysis technique (Hyndman and Athanasopoulos 2018). Given that this is a heating load profile, it is also expected to be highly correlated with weather conditions. Considering these in mind, we conducted EDA based on time components and weather conditions as shown in Figure 2.

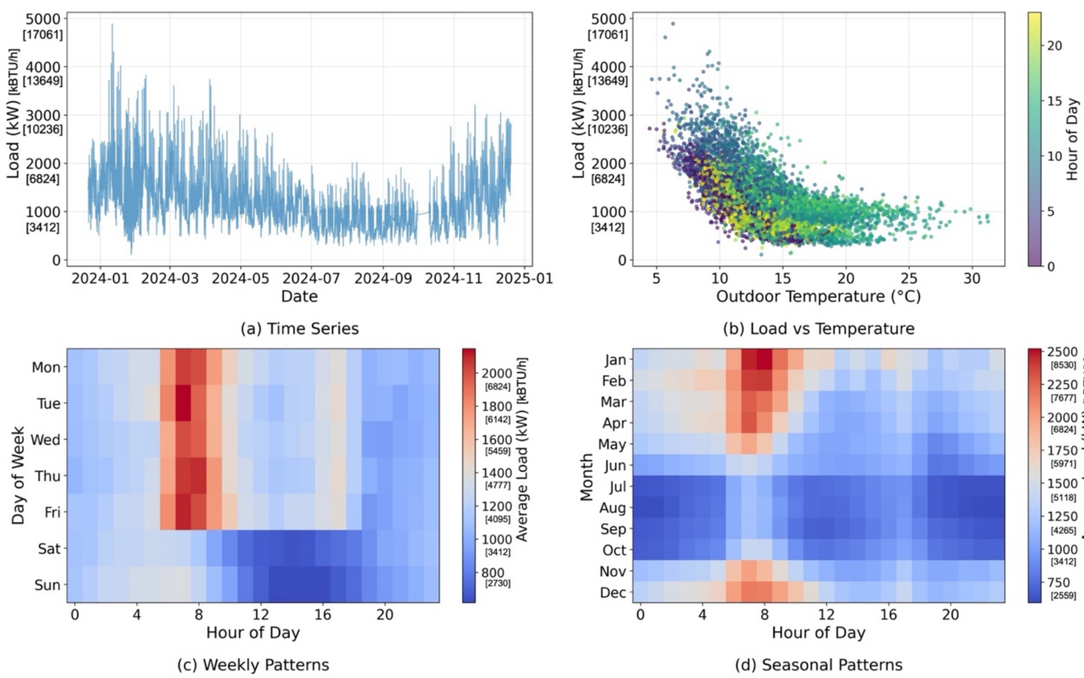

Figure 2 Exploratory data analysis of the annual heating load profile, showing (a) the time-series data for one year, (b) the correlation between load and outdoor temperature, colored by the hour of day, (c) the average weekly load pattern, and (d) the average seasonal load pattern.

The raw time-series data for one year of operation is shown in Figure 2(a), which illustrates the significant fluctuations in the building's heating demand. It is worth noting that a period of missing data occurs in October. Figure 2(b) explores the relationship between the heating load and the primary external driver, outdoor air temperature. As expected, a strong negative correlation is evident: the heating load increases as the outdoor temperature decreases. However, while the relationship with temperature is clear, the impact of the hour of day (represented by the color scale) is difficult to discern from this plot alone, suggesting that other temporal patterns influence the load. To better visualize these temporal characteristics, the data was aggregated into weekly and seasonal heatmaps. Figure 2(c) reveals a distinct weekly pattern, with weekdays (Monday-Friday) showing a different and higher load profile compared to weekends (Saturday-Sunday). This indicates a strong dependence on the building's operational schedule and occupancy. The seasonal patterns are shown in Figure 2(d), where the load is highest in the winter months and decreases to a minimum during the summer. The gradual increase and decrease in heating demand through the shoulder seasons reinforce the observation that the load profile is more closely related to continuous weather conditions than to discrete seasons.

To further investigate the combined effects of these variables, the annual load data was disaggregated by the day of the week, as shown in Figure 3. This detailed visualization confirms the patterns identified in the initial EDA and reveals the strong interplay between time-based schedules and weather conditions. The profiles for weekdays (Monday-Friday) are distinctly different from those for the weekend (Saturday-Sunday), with weekdays exhibiting a higher average load and a more pronounced daily shape driven by the building's operational schedule. Within each day's plot, the effect of outdoor air temperature is clearly visible through the color mapping. Colder days, represented by purple and blue lines, consistently correspond to higher heating loads, while warmer days, represented by yellow and green lines, result in lower loads. This confirms that outdoor temperature is a primary driver of the heating demand's magnitude, regardless of the day of the week or time of day. The analysis of these daily profiles reinforces the conclusion that the heating load is a complex function of the day of the week, the hour of the day, and outdoor air temperature, highlighting the need for a model that can capture these interacting features.

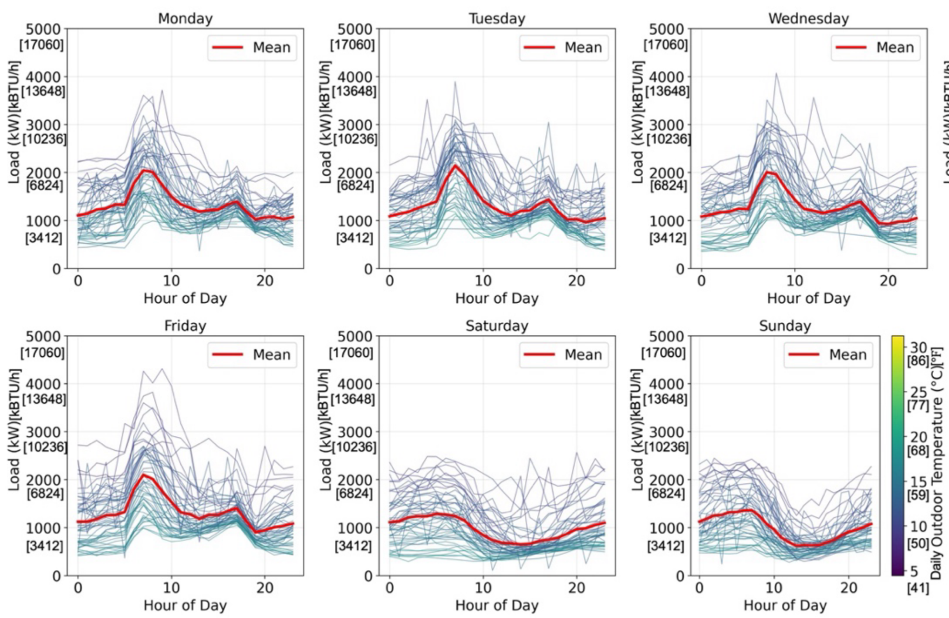

Figure 3 Daily heating load profiles for each day of the week. Each plot shows individual daily load profiles (thin lines) and the mean profile (red line). The color of the individual profiles corresponds to the daily average outdoor air temperature, as indicated by the color bar.

Model structure

Based on the exploratory data analysis, the building's heating load is driven by a combination of deterministic, time-based patterns and external weather conditions, along with inherent stochastic variability. To capture these dynamics accurately, a modeling approach is required that can learn the underlying structure from known features while remaining flexible enough to represent the unexplained, random component of the load. Therefore, we use a Gaussian Process (GP) model, a non-parametric, probabilistic approach well-suited for this task. The GP model's structure is defined by its input features and a covariance function (kernel), which describes the relationship between data points.

The model is designed to predict the heating load based on an input vector, x, which contains the features identified during the data analysis: $x = [t_{\mathrm{h}}, t_{\mathrm{d}}, T_{\mathrm{oa}}]$, where t_{h} is the hour of the day (0-23), t_{d} is the day type (1 for weekday, 0 for weekend), and T_{oa} is the outdoor air temperature. The core of the GP model is its covariance function, or kernel $k(x, x')$, which defines the relationship between any two input points. To model the time-series nature of the heating load, we selected an additive GP model structure, a technique demonstrated for decomposing time-series data in (Gelman 2014). From the additive time-series GP model in (Gelman 2014), we use a Standard Periodic kernel for weekly repetitive patterns by differentiating the day type with a Radial Basis Function (RBF) kernel (k_{weekly}). In addition, we use the Rational Quadratic kernel to account for the impact of outdoor air temperature ($k_{T_{\mathrm{oa}}}$). We also tested the GP model that has the interaction effect of k_{weekly} and $k_{T_{\mathrm{oa}}}$, but we didn't observe any noticeable differences in terms of accuracy.

The full additive kernel is expressed as:

$$k(x, x') = \underbrace{\left[\sigma_1^2 \exp\left(-\frac{2\sin^2(\pi|t_h - t_h'|/P)}{l_1^2}\right)\exp\left(-\frac{(t_{\mathrm{d}} - t_{\mathrm{d}}')^2}{2l_2^2}\right)\right]}_{k_{\mathrm{weekly}}} + \underbrace{\left[\sigma_2^2\left(1 + \frac{(T_{oa} - T_{\mathrm{oa}}')^2}{2\alpha l_3^2}\right)^{-\alpha}\right]}_{k_{T_{\mathrm{oa}}}} \tag{1}$$

where the first component, k_{weekly}, captures the recurring daily and weekly patterns. It is constructed as a product of a Standard Periodic kernel, which models the 24-hour daily cycle ($P = 24$), and a RBF kernel that acts as a switch to differentiate between weekday and weekend patterns. In this formulation, the hyperparameter σ_1^2 controls the amplitude of the daily pattern, l_1 controls how smoothly the pattern changes over the hours, and l_2 controls the similarity between weekday and weekend patterns. The second component, $k_{T_{\mathrm{oa}}}$, models the impact of outdoor air temperature on the heating load using a Rational Quadratic kernel, which is capable of capturing effects across multiple temperature scales. Here, the hyperparameter σ_2^2 controls the overall magnitude of the temperature's impact, l_3 defines the temperature range over which the heating loads are correlated, and α is a scale-mixture parameter that allows the model to capture both large- and small-scale temperature variations.

We implemented the Gaussian Process model using the GPy library (GPy 2012), a popular open-source framework for GP modeling in Python. To determine the optimal values for the kernel's hyperparameters (i.e., $\sigma_1^2, l_1, \sigma_2^2, l_2, \alpha, l_3$), the model was trained on a three-month segment (April-June) of the measured heating load data. This period was selected because it captures a sufficient range of outdoor air temperatures. The training process involved optimizing the hyperparameters to maximize the marginal likelihood of the observed data, with the remainder of the annual dataset reserved for testing the model's predictive performance.

Although the entire dataset could have been used for training—since the primary goal is to generate typical load profiles—we intentionally used a train/test split to diagnose potential overfitting. A GP model trained on the entire dataset can easily overfit by acting as a high-order interpolator, effectively memorizing the data rather than learning the underlying relationship between the features ($[t_{\mathrm{h}}, t_{\mathrm{d}}, T_{\mathrm{oa}}]$) and the heating load. In such a scenario, the model would incorporate an unnecessarily large stochastic component, failing to attribute variations to the deterministic features. By splitting the data, we can evaluate the model's performance on the unseen test set to validate its ability to generalize and ensure it has captured the true effect of the input features.

RESULTS

This section presents the performance of the trained GP model. First, we evaluate its predictive accuracy and uncertainty quantification on both the training and testing datasets. Second, we interpret the physical meaning of the optimized kernel hyperparameters.

Prediction results on training/testing data

The predictive performance of the GP model is illustrated in Figure 4, which displays the model's output against the actual measured heating load for representative weeks from both the training and testing periods. As shown in

Eq. (1), once the model is trained, we can predict the load profile based on feature values such as hour of day, day type, and outdoor air temperature. In other words, for days with the same feature values and sequences, the prediction results are identical. However, the GP model accounts for uncertainties arising from both measured disturbances (e.g., outdoor air temperature) and unmeasured disturbances (e.g., occupancy profiles) in its predictive uncertainty. By sampling load profiles for given feature values, we can generate profiles for different scenarios, such as varying occupancy schedules. The top row of plots shows the model's fit to the training data (April, May, and June). The GP predicted mean (blue line) closely tracks the actual observed data (black line), indicating that the model has effectively captured the primary patterns and dynamics present in the training set. Critically, the 95% confidence interval (shaded blue area) successfully encompasses the vast majority of the true measurements, demonstrating a well-calibrated representation of model uncertainty. The bottom row of plots validates the model's performance on unseen testing data (January, September, and December). The model maintains its strong predictive accuracy, with the predicted mean (red line) continuing to follow the daily and weekly fluctuations of the actual heating load. The 95% confidence interval (shaded red area) also remains robust, containing most of the observed data points. Furthermore, we present CV(RMSE) and NMBE, which are widely accepted in retrofit analysis (CV(RMSE<30%, NMBE<15% (Haberl, Claridge, and Culp 2005)), as references. The consistent performance across both datasets confirms that the GP model has learned the underlying relationship between the input features and the heating load, rather than simply overfitting to the training data. The model successfully captures the deterministic effects of these features while also providing a reliable quantification of uncertainty.

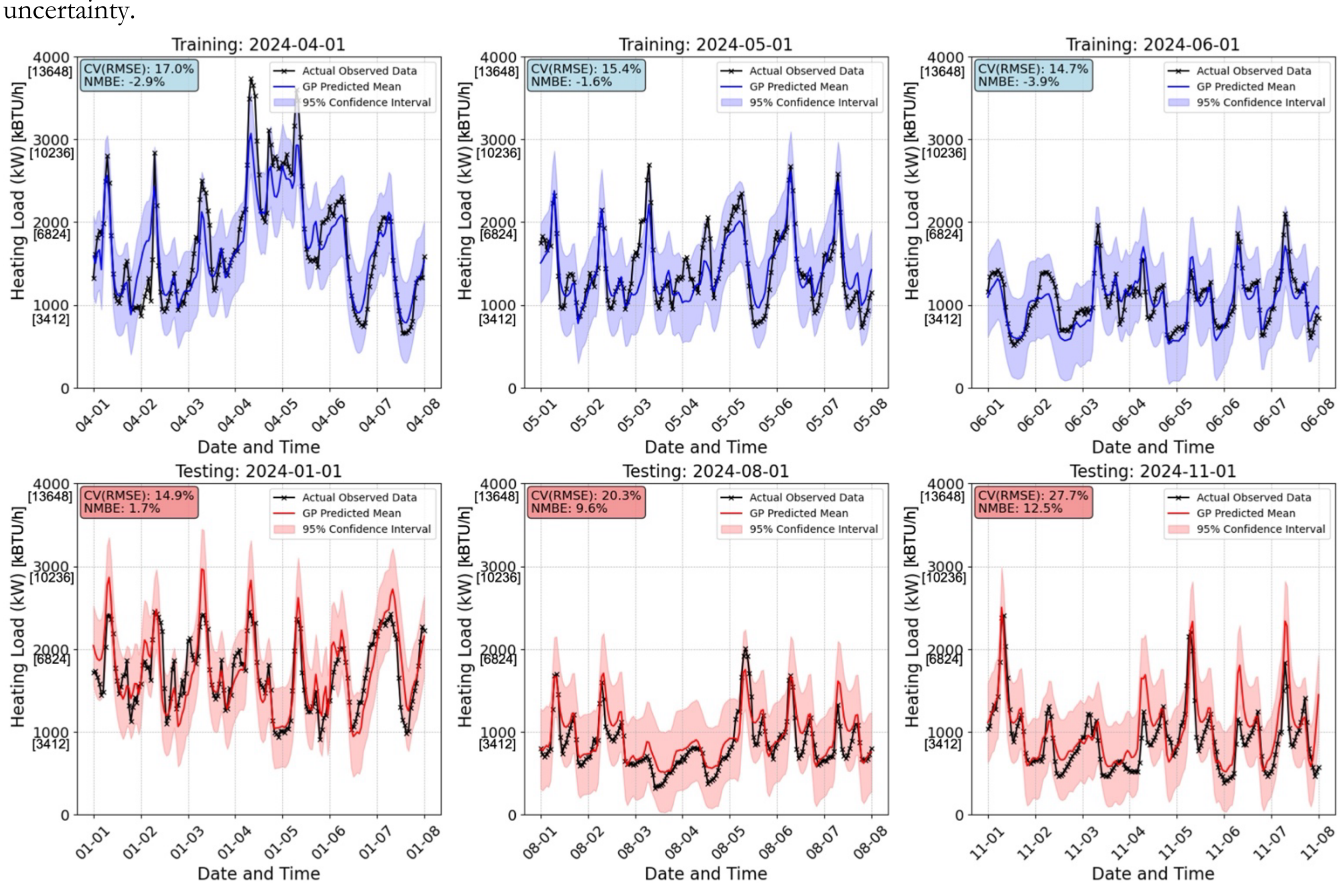

Figure 4 GP Model prediction performance on Training (Blue) and Testing (Red) data with predictive uncertainties.

Interpretation of GP kernels

Decomposing the GP model's prediction into its kernel contributions provides physical insight into the building's heating load, as visualized in Figure 5. The weekly pattern kernel (middle plot) successfully isolates the building's operational schedule. Its contribution is high during occupied weekday hours and minimal during nights and weekends, which accurately reflects the influence of the building's schedule. The outdoor air temperature kernel (bottom plot) captures the weather-driven heating load. Its contribution is inversely correlated with the outdoor temperature (dashed line), increasing as temperatures fall and decreasing as they rise, clearly isolating the physical impact of weather on the building's heating demand.

CONCLUSIONS AND FUTURE WORK

In conclusion, this study successfully developed and validated a probabilistic, load-based framework using a GP model to support HVAC retrofit decisions. The model demonstrated strong predictive accuracy on unseen data, and its kernel decomposition provided clear, interpretable insights into how building schedules and weather impact heating loads. Building on this work, future efforts will focus on expanding the framework to include a corresponding cooling load model for year-round analysis and researching more computationally efficient methods for uncertainty quantification to overcome the expense of Monte Carlo simulations.

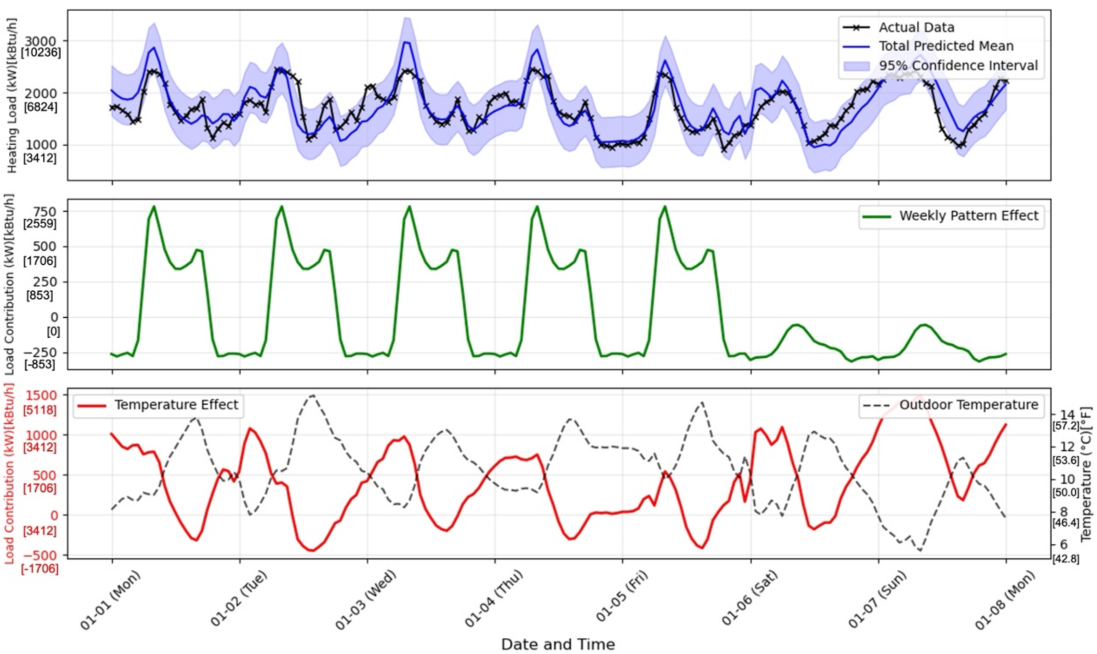

Figure 5 A visualization of the GP model's total prediction (top) and its decomposition into the weekly pattern (middle) and outdoor temperature (bottom) kernel contributions.

ACKNOWLEDGMENTS

This manuscript has been authored by an author at Lawrence Berkeley National Laboratory under Contract No. DE-AC02-05CH11231 with the U.S. Department of Energy. This work was supported by the California Energy Commission under grant number EPC-23-304. The U.S. Government retains, and the publisher, by accepting the article for publication, acknowledges, that the U.S. Government retains a non-exclusive, paid-up, irrevocable, world-wide license to publish or reproduce the published form of this manuscript, or allow others to do so, for U.S. Government purposes.

REFERENCES

Bai, Xi, Qi Tang, Jianing Luo, Yubo Mao, Caihua Liang, and Xiaosong Zhang. 2024. "Optimizing Energy Efficiency in Multi-Chiller Systems: A Comprehensive Modelica-Based Approach." *Journal of Building Engineering* 95 (110087): 110087.

Bhattacharya, Arnab, Soumya Vasisht, Veronica Adetola, Sen Huang, Himanshu Sharma, and Draguna L. Vrabie. 2021. "Control Co-Design of Commercial Building Chiller Plant Using Bayesian Optimization." *Energy and Buildings* 246 (111077): 111077.

California Energy Design Assistance. 2023. "Benefits of Heat Recovery Chillers." https://californiaeda.com/benefits-of-heat-recovery-chillers-as-they-contribute-to-decarbonization-and-building-electrification.

Carrier. 2022. "Heat Recovery from Chilled Water Systems: Applications for Heat Reclaim Chillers." 04-581025–01.

Catrini, P., D. Panno, F. Cardona, and A. Piacentino. 2020. "Characterization of Cooling Loads in the Wine Industry and Novel Seasonal Indicator for Reliable Assessment of Energy Saving through Retrofit of Chillers." *Applied Energy* 266 (114856): 114856.

Chen, Jianli, Rajendra Adhikari, Eric Wilson, Joseph Robertson, Anthony Fontanini, Ben Polly, and Opeoluwa Olawale. 2022. "Stochastic Simulation of Occupant-Driven Energy Use in a Bottom-up Residential Building Stock Model." *Applied Energy* 325 (119890): 119890.

Gelman, Andrew. 2014. *Bayesian Data Analysis*. Third edition. Chapman & Hall/CRC Texts in Statistical Science. Boca Raton: CRC Press.

Gill, Brandon. 2021. "Solving the Large Building All-Electric Heating Problem." *ASHRAE Journal* 63 (10).

GPy. 2012. "GPy: A Gaussian Process Framework in Python." 2012. http://github.com/SheffieldML/GPy.

Haberl, J., D. Claridge, and C. Culp. 2005. "ASHRAE's Guideline 14-2002 for Measurement of Energy and Demand Savings: How to Determine What Was Really Saved by the Retrofit." https://oaktrust.library.tamu.edu/bitstreams/132afe4f-7ca7-448c-a818-f7fcba08f3cd/download.

Heemer, Jim, Alexander Mitrovic, and M. Scheer. 2011. "Increasing Central Plant Efficiency via a Water to Water Heat Pump." *Pharmaceutical Engineering*.

Hinkelman, Kathryn, Jing Wang, Wangda Zuo, Antoine Gautier, Michael Wetter, Chengliang Fan, and Nicholas Long. 2022. "Modelica-Based Modeling and Simulation of District Cooling Systems: A Case Study." *Applied Energy* 311 (118654): 118654.

Huang, Sen, Wangda Zuo, and Michael D. Sohn. 2016. "Amelioration of the Cooling Load Based Chiller Sequencing Control." *Applied Energy* 168 (April): 204–15.

Hyndman, R. J., and G. Athanasopoulos. 2018. "Forecasting: Principles and Practice." https://www.academia.edu/download/64659947/Athanasopoulos__George__Hyndman__Rob_J._-_Forecasting__Principles_and_Practice_2018.pdf.

Lee, Sang Hoon, Tianzhen Hong, Mary Ann Piette, and Sarah C. Taylor-Lange. 2015. "Energy Retrofit Analysis Toolkits for Commercial Buildings: A Review." *Energy (Oxford, England)* 89 (September): 1087–1100.

Liang, Chenjiyu, Yanqing Lin, Cuiling Wang, Wentao Wang, Baolong Wang, Xianting Li, Jun Cheng, and Yuan Wang. 2025. "Stepwise Commissioning Method to Predict Energy Efficiency Improvements in Chiller Plant via Retrofitting." *Journal of Building Engineering* 105 (112586): 112586.

Tommasi, Luciano De, Hassan Ridouane, Georgios Giannakis, Kyriakos Katsigarakis, Georgios N. Lilis, and Dimitrios Rovas. 2018. "Model-Based Comparative Evaluation of Building and District Control-Oriented Energy Retrofit Scenarios." *Buildings* 8 (7): 91.

Wang, Zihao, Wenkai Zhang, Hang Fan, Chaobo Zhang, Yang Zhao, and Zemao Huang. 2023. "An Uncertainty-Tolerant Robust Distributed Control Strategy for Building Cooling Water Systems Considering Measurement Uncertainties." *Journal of Building Engineering* 76 (107162): 107162.

Wetter, Michael, Wangda Zuo, Thierry S. Nouidui, and Xiufeng Pang. 2014. "Modelica Buildings Library." *Journal of Building Performance Simulation* 7 (4): 253–70.

Wilson, Eric J. H., Andrew Parker, Anthony Fontanini, Elaina Present, Janet L. Reyna, Rajendra Adhikari, Carlo Bianchi, et al. 2022. "End-Use Load Profiles for the U.S. Building Stock: Methodology and Results of Model Calibration, Validation, and Uncertainty Quantification." United States: National Renewable Energy Lab.(NREL), Golden, CO (United States). https://doi.org/10.2172/1854582.