# Scalable Identification and Control of Residential Heat Pumps: A Minimal Hardware Approach

Zachary E. Lee, K. Max Zhang*

*Sibley School of Mechanical and Aerospace Engineering, Cornell University, Ithaca, NY 14853, USA*

**Abstract**

While model predictive control is a widely studied tool for integrating residential heat pumps with the smart grid, modeling difficulties and hardware requirements are significant barriers to adoption. In response, we present a nonintrusive plug-and-play methodology for model identification and model predictive control using only two increasingly popular smart-home devices: a smart thermostat and smart electricity meter, which makes the overall approach highly scalable. However, this minimal hardware approach requires new methods to overcome modeling challenges due to the lack of data diversity. A data-driven heat pump power model is identified without the need for submetering by using energy disaggregation. By incorporating the heat pump control data from the smart thermostat, we provide implicit load classification data to the model, substantially simplifying the signal decomposition task. Building and heat pump model parameters are then identified by a novel algorithm that combats overfitting caused by limited state excitation. Finally, we show the advantage of this scalable approach in an aggregate model predictive controller enabled by the communication ability provided by the smart thermostat. Results show that this aggregate control approach can provide improved grid services, such as reduced peak demand, while at the same time reducing energy consumption and improving thermal comfort. The proposed method has the potential to greatly lower the barrier for residential energy aggregation.

*Keywords:*
Grid Interactive Buildings; System Identification; Machine Learning; Model Predictive Control; Building Climate Control

## Highlights

- A minimal hardware approach to heat pump and smart-grid interactivity is presented.
- Only requires a smart thermostat and electricity meter and is thus highly scalable.
- A novel system identification method combats overfitting from low system excitation.
- The scalable approach allows for improved grid services through aggregate control.
- Providing aggregate grid services does not reduce local controller performance.

*Corresponding author: kz33@cornell.edu

## 1. Introduction

Building electrification is projected to see a rapid rise in the coming years due to the transition to a carbon-neutral energy system. Over 60% of U.S. homes (and 80% of homes in the Northeastern US) use on-site fossil fuel-based space heating, contributing to 10% of the nation's greenhouse gas emissions [1]. Recent advances in heat pump technology have lowered capital costs and increased their efficiency [2], particularly in cold-climates [3], making them one of the preferred alternatives to fossil fuel-based home heating in clean energy initiatives around the world [4]. Since heat pumps can provide heating in the winter and cooling in the summer, they displace both fossil fuel heating and inefficient window air conditioners used for cooling. However, these new electrical loads can have adverse effects on an increasingly renewable-based electrical grid, where the new demand profiles and peak loads may not follow the availability of renewable resources.

In contrast to fossil fuel-based heating, the temporal profile of heat pump operation throughout the day significantly affects their efficiency and societal cost. Heat pump efficiency, or coefficient of performance (COP), is inversely proportional to the difference between indoor and outdoor temperatures, causing the efficiency to vary on the order of $\pm 20\%$ throughout the day. In addition, while retail electricity prices often remain relatively constant, wholesale electricity prices can vary significantly throughout the day depending on electrical demand and the availability of renewable energy. Without sufficient grid electricity storage or demand flexibility, satisfying these new heat pump electrical loads in times of high wholesale electricity prices could raise retail electricity prices, discouraging heat pump adoption and disproportionately affecting low-income families, whose energy bills can already encompass between 9 and 19% of their annual income [5]. These effects provide a large incentive for optimally controlling heat pumps, both to reduce operating costs and to facilitate the integration of renewable heating and cooling options.

Smart control of heat pumps for increased efficiency and better integration with the electrical grid generally takes two primary approaches: local and aggregate control [6]. Local control focuses on increasing efficiency and reducing costs by modeling and controlling each residence independently, either through detailed and specific physical modeling or through data collected by sensors installed throughout the building [7]. Model predictive control (MPC) is a widely popular technique for local heat pump control and has been used to respond to dynamic electricity prices [8] and to reduce emissions [9]. In Ref. [10], a local MPC was included in a hierarchical controller to respond to demand response signals sent by the higher level controller. However, each local controller's optimal control sequence cannot consider the control behavior of any other heat pump, and therefore cannot coordinate with each other to increase their aggregate demand response potential.

Aggregate control primarily seeks to provide a service to the electrical grid, such as flexibility for demand response [11] or even more complex services such as managing energy imbalance [12]. Rather than using detailed thermodynamic models of each home, they model the multi-unit aggregate dynamics and use techniques such as global setpoint modifications [13] or priority list controllers [14] to control the aggregate power to follow a desired objective. These controllers are often designed under the assumption that there is minimal hardware available in each home, and therefore bi-directional communication is not easily available. While these studies can effectively shape the aggregate load profile to provide certain grid services, they do not attempt to increase the efficiency, thermal comfort, or operating cost of each individual residence. Moreover, Ref. [15] shows that modeling complexity, including each heat pump's performance and their minimum cycle times, has a significant effect on an aggregate controller's key performance indicators.

Recently, smart-home devices, which contain built-in communication hardware, are becoming increasingly ubiquitous, opening up a widespread avenue for implementing smart control of heat pumps. For example, smart thermostats have seen rapid adoption in recent years, and are projected

to be in 40% of US homes by 2021 [16]. Smart thermostats provide users with a wide range of energy saving tools, including custom setpoint schedules, varying comfort modes such as "sleep" and "away", and even occupancy sensors that reduce energy consumption when the home is empty. Smart thermostat control complexity is still relatively limited, though, and predictive control and building modeling are rarely used to increase heat pump efficiency or better respond to grid signals like dynamic electricity prices or demand response calls. However, some thermostat manufacturers have begun to release APIs to remotely connect to the thermostat, allowing the potential for more complex control algorithms to be implemented from a centralized location.

Similarly, smart electricity meters capable of recording a home's electricity consumption over relatively short time intervals are also being rapidly installed by utilities. Some utilities have developed dashboards allowing users to interact with their data. However, while the information from smart meters is limited to whole-home[1] power, energy disaggregation has emerged as an experimental technique to estimate device-level power based on the whole-home power profile. This data-driven technique can be used to provide real-time energy feedback for users to reduce their energy consumption.

Adoption of these smart-home devices has begun to shape user behavior, and as these devices become more widespread, the consequences of these behavioral changes will become more significant. While the effects of energy consumption feedback from smart-meters has been widely studied [17], the effect of smart thermostats has only recently been explored [18], facilitated by the newly available Ecobee Donate your Data smart thermostat dataset [19]. From an analysis of this dataset, we find that over 90% of smart thermostat users use custom setpoint schedules, and the combined outcome of these setpoint schedules can have a significant impact on the multi-unit aggregate demand profile. Though smart thermostats do reduce energy consumption on a local level, the current local control method based on setpoint schedules can potentially introduce unintended consequences for the aggregate energy system by creating new or more pronounced peaks in aggregate demand. Therefore, rather than constant setpoints used in other aggregate heat pump control studies [20], an aggregate controller must consider realistic setpoint schedules to mitigate the potential negative consequences that smart thermostat setpoint schedules can have on the overall energy system.

This paper presents a method that requires minimal hardware to control smart thermostats to further reduce local energy costs, while at the same time mitigating their aggregate effect on the energy system. The key benefit of this approach is that it is a plug-and-play solution for residential heat pump and smart grid interactivity that requires only two increasingly common smart-home devices: a smart thermostat and a smart electricity meter. By limiting ourselves to only the data and communication provided by these devices, this methodology is highly scalable, non-intrusive, and easy to implement. We show that the scalability of this approach can not only enable widespread implementation of more efficient heat pump control, but can also provide the ability for centralized control of heat pump aggregations, which improves our ability to provide grid services such as peak load management. The residential energy aggregation market has been long recognized as a major challenge as the hardware and labor costs to install building-level and appliance-level load monitoring devices are usually not justifiable economically for small loads. As smart thermostats and utilities' smart electricity meters have become increasingly ubiquitous, the proposed method has the potential to greatly lower the barrier for residential energy aggregation.

While this minimal hardware approach provides significantly reduced barriers to adoption, the main technical challenge is managing the modeling difficulties associated with the limited diversity

[1]Instead of the term "aggregate power" more commonly seen in energy disaggregation literature, we use the term "whole-home power" to avoid confusion with the multi-unit aggregate control.

of data available.

First, data is collected from the thermostat and electricity meter in order to identify a data-driven model of the heat pump power, which can vary significantly from published manufacturer data due to control delays, installation differences, and transient effects. We therefore use energy disaggregation to estimate the heat pump's power based on the whole-home power profile. We find that incorporating the heat pump control data from the smart thermostat substantially improves the ability to predict the heat pump power using machine learning techniques. We then use a linear autoregressive data-driven model to identify the specific heat pump's power dynamics.

Second, the building thermodynamics must be identified based on the limited amount of data obtained from the two devices. Due to insufficient excitation of the building thermodynamics and limited sensor coverage, typical thermal circuit identification and black-box modeling methods significantly overfit and can be unreliable. We show that while popular black-box methods can accurately predict temperature under the existing thermostat control scheme, their open-loop dynamics do not follow thermodynamic principles and consequently do not perform well in model-based control. Instead, we adopt the common two-state grey-box thermal circuit model but introduce an improved algorithm for identifying optimal model parameters when limited input data is available.

In order to show the advantage of this approach's scalability, we propose a novel MPC approach for energy aggregation, referred to as aggregate MPC in this paper, that controls a community of smart thermostats for improved aggregated energy management. Since we have access to each building's local thermodynamics through our scalable system identification method, we can shape the aggregate electrical demand profile for certain aggregate grid services while still maintaining each building's thermal comfort and energy efficiency. This approach differs from other aggregate control methods which often do not consider local energy efficiency and must artificially change residents' setpoints. Thus, we combine the energy efficiency and thermal comfort benefits of local MPC control with the system level benefits of aggregate control.

This approach also provides the ability to incorporate two improvements to the MPC formulation that are often neglected in literature: minimum heat pump on and off times to prevent shortcycling and actual smart thermostat setpoint schedules from real occupants. These improvements are vital for understanding real-world limitations, particularly the effect that real smart thermostats setpoint schedules can affect peak load and the aggregate electrical consumption profiles. Finally, we illustrate the performance of this approach by minimizing the aggregate peak load of a population of residences. We show that our aggregate controller can significantly lower the aggregate peak load without a significant increase in energy consumption or decrease in thermal comfort when compared to a local control approach.

The remainder of the paper is organized as follows. Section 2 outlines how each of the system models are identified from data, including heat pump power, building thermodynamics, and disturbance predictions. Section 3 formulates the model predictive control problem. Section 4 discusses the control simulation configuration and results. Section 5 concludes the paper.

## 2. Scalable System Identification

In this section we present a general methodology for data-driven identification of a residence's heat pump and building thermodynamic models assuming that each residence is equipped with at least one smart thermostat and one smart electricity meter. Data measured by the thermostat include indoor air temperature and relative humidity, thermostat setpoint, boolean occupancy, and heat pump runtime in seconds. The smart electricity meter provides whole-home power, which includes the contributions from both the heat pump power $P_{\mathrm{HP}}$ and the other miscellaneous power

consumption from the occupant that we define as indoor loads $P_{\text{in}}$. While some weather data, such as outdoor temperature can be provided through the smart thermostat, other weather data such as solar irradiation and humidity can be readily acquired from a nearby weather station or forecasting service. While we analyze the system in heating, this procedure can be equally applied to cooling as well. The remainder of this section presents how each of the models is developed and incorporated into the methodology, which is outlined in Fig. 1.

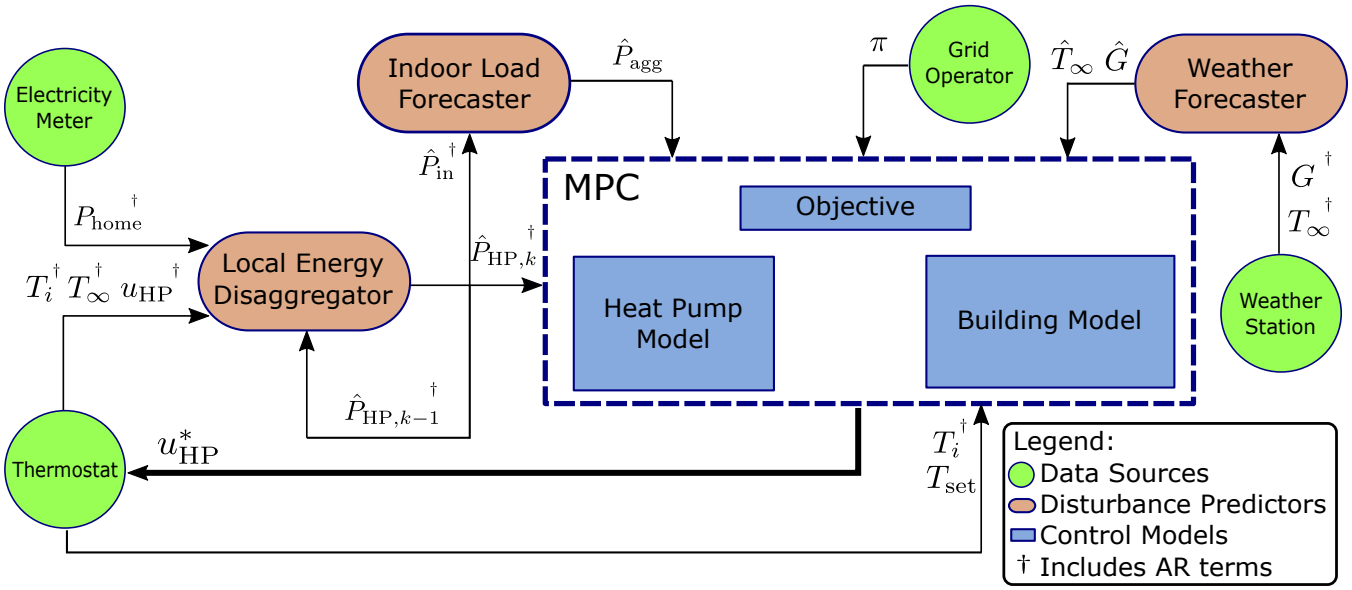


Figure 1: Methodology outline: Each data source provides the data required for training the disturbance and control models and for running the model predictive controller. Some data inputs include autoregressive (AR) terms to aid in model prediction. Forecasts from the prediction models provide disturbance predictions over the MPC horizon.

### *2.1. Test System Description*

To test our methodology, we use data from an occupied residential building representative of a majority of heat pump-heated US homes, a sample of which is shown in Fig. 2. The residence, located in Ithaca, NY, is a well insulated 130 m$^2$ building heated by a single-stage air to air split system heat pump (model number TH4B2421SA). A single ecobee4 smart thermostat controls the heat pump and records data every 5 minutes. The thermostat is set to its default setpoint settings, with a setpoint of 21°C during the "home" mode from 6:30 am to 11:30 pm, and a setpoint of 18°C during "sleep" mode from 11:30 pm to 6:30 am. Additionally, the thermostat will periodically set back the setpoint when no occupancy is detected for a certain amount of time. The heat pump, air handler, and total building power consumption are all measured through current transformers located in the circuit panel. Weather data, including temperature, solar irradiation, dewpoint, wind speed, and pressure, were obtained from the NREL National Solar Radiation Database [21]. Data were collected for this building for the entire month of November 2019, resulting in 7000 5-minute timesteps.

### *2.2. Heat Pump Power Modeling using Energy Disaggregation*

A majority of heat pump control studies model heat pump power and performance either through steady-state physics-based models or from manufacturer's lab data [8, 22]. Instead, we opt for a purely data-driven approach due to various effects that arise in real-world operation that cause significant deviations from these assumptions. First, the heat pump has a warm-up time, taking

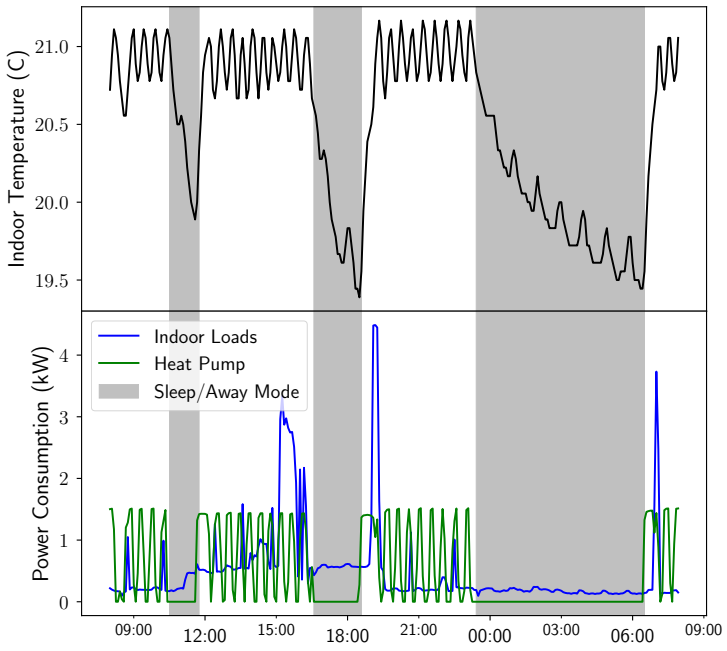


Figure 2: Data for a typical day in November from our test system. Indoor loads include any non-heat pump related power consumption.

up to 10 minutes to reach its steady-state power after being turned on. Second, there can be a significant control delay of up to two minutes between the thermostat calling for heat and the heat pump actually switching on, due to varying internal heat pump or thermostat control mechanisms. Combined, these two factors can cause the average power over a timestep where the heat pump is turning on to be significantly lower than the steady-state models predict. In addition, due to installation differences in the field, the actual efficiency and power can vary widely from the manufacturer lab data [23]. Because of these issues, we use an autoregressive (AR) data-driven model that can provide a more accurate power prediction by taking into account the combined effects of control delays, warm-up times, and differences between field and lab operation.

A key limitation preventing this data-driven AR approach so far has been the requirement to submeter the heat pump's electrical consumption in order to obtain training data, which can be costly and require a professional to install. We solve this problem by using energy disaggregation to predict the heat pump power profile based on the whole-home power profile measured by the smart electricity meter, eliminating the need for a heat pump submeter in field operation. Energy disaggregation, sometimes called non-intrusive load monitoring, is a method for estimating device-level (disaggregated) power consumption based on the whole-home (aggregate) power profile. It can be classified into two main approaches: load classification and signal decomposition. Load classification is used to determine which devices are on at any point in time and is accomplished using methods like event detection [24] and hidden Markov models [25]. Signal decomposition aims to separate the individual power signals from each device from the whole-home power signal, and methods for this task include sparse coding [26], change detection [27], and Sum-to-$k$ constrained Non-negative Matrix Factorization [28]. Because these papers attempt the relatively difficult task of

estimating power information for many different devices based on the whole-home power signature, they often require significant amounts of training data and complex models. In contrast, our method includes the heat pump control data provided the smart thermostat, which provides implicit load classification information and significantly simplifies the signal decomposition task.

To accomplish the disaggregation, we use random forest regression. Random forest is a supervised learning technique that creates an ensemble of different decision trees and outputs the mean of each decision tree's individual prediction. By taking the mean of each tree's output, random forest combats the overfitting that commonly occurs in individual decision trees. We find that random forest regression outperforms other supervised learning regression methods for this task due to its ability to efficiently incorporate the load classification information available from the control signal into its decision trees. Shown by the Local Energy Disaggregator block in Fig. 1, input features require (1) the model's output for the previous timestep $P_{\mathrm{HP},k-1}$, (2) the heat pump control signal $u_{\mathrm{HP}}$, (3) the whole-home power $P_{\mathrm{home}}$, (4) indoor temperature $T_i$, and (5) outdoor temperature $T_\infty$, each of which contain AR terms from the previous timestep. The output is the disaggregated heat pump power $\hat{P}_{\mathrm{HP},k}$ for the current timestep. A prediction of other indoor loads $\hat{P}_{\mathrm{in}}$ can then be found by subtracting the heat pump power from the total power at that timestep. We use 6000 datapoints for training, with 500 datapoints used for five-fold cross-validation to determine the optimal hyperparameters using the Sci-kit Learn Python package [29]. The remaining samples were used for testing. Disaggregation test performance is shown in Fig. 3 with the final test RMSE at 0.037 kW, equivalent to around 2% of the steady-state heat pump power.

We then use the disaggregated heat pump power to train a linear AR model of the heat pump power that can be used in linear MPC. In this case, linear regression determines the optimal coefficients $\boldsymbol{\theta}$ to minimize the residual sum of squares between the predicted heat pump power $\hat{P}_{\mathrm{HP},k}$ and the actual power $P_{\mathrm{HP},k}$ at the current timestep $k$. While the disaggregated power is used as labels for training the model, test performance is evaluated based on the difference between the prediction and the submeter-measured heat pump power. As input features $\boldsymbol{x_k}$, we include the percentage of time the thermostat was calling for heat over the timestep $u_{\mathrm{HP},k}$, indoor temperature $T_i$, and outdoor temperature $T_\infty$. To capture the effects of warm-up time and control delays, we also include AR terms $\boldsymbol{x}_{k-1}$ to $\boldsymbol{x}_{k-p}$ and $P_{\mathrm{HP},k-1}$ to $P_{\mathrm{HP},k-p}$, where $p$ is the number of AR timesteps. This model is given as:

$$\begin{aligned} &\min_{\theta} \sum_{k+p}^{K} (P_{\mathrm{HP},k} - \hat{P}_{\mathrm{HP},k})^2 \\ &\text{subject to } \hat{P}_{\mathrm{HP},k} = \boldsymbol{\theta}_x \begin{bmatrix} \boldsymbol{x}_k \\ \boldsymbol{x}_{k-1} \\ \dots \\ \boldsymbol{x}_{k-p} \end{bmatrix} + \boldsymbol{\theta}_P \begin{bmatrix} P_{\mathrm{HP},k-1} \\ \dots \\ P_{\mathrm{HP},k-p} \end{bmatrix} \end{aligned} \tag{1}$$

The 7000 timesteps of data samples were divided into 6000 timesteps for training, and 500 timesteps each for validation and testing. Since the heat pump's duty cycle changes based on the heat load, we divide the data into four hour segments and shuffle it such that the training, validation, and test sets contain the same distribution of heat pump operation. We use five-fold cross validation to determine that one AR term is sufficient to capture the transient effects.

We show the improvement of our data-driven AR model by comparing it to the more conventional steady-state (SS) models (no AR terms) where parameters are based on our field experimental data and where parameters are determined from the manufacturer performance data. Fig. 4 shows the model prediction results for the data-driven AR model, the field data-driven SS model, and

the SS model derived from the manufacturer performance data [30] using the disaggregated heat pump power as training labels. One interesting result is that in this case disaggregation does not significantly affect model performance. The disaggregation test error tends to follow a Gaussian distribution, and therefore over many datapoints, the Gaussian noise introduced by the disaggregation into the heat pump power training labels does not adversely affect the resulting heat pump power model. If the disaggregation test error had more significant bias, such as consistent over- or underpredicting, then the heat pump model trained on disaggregated data would be negatively effected. The large difference between the manufacturer data-based model and the field data-based models is likely due to field installation settings that differ from the performance datasheet, underscoring the need for field-data based models that can capture this variability.

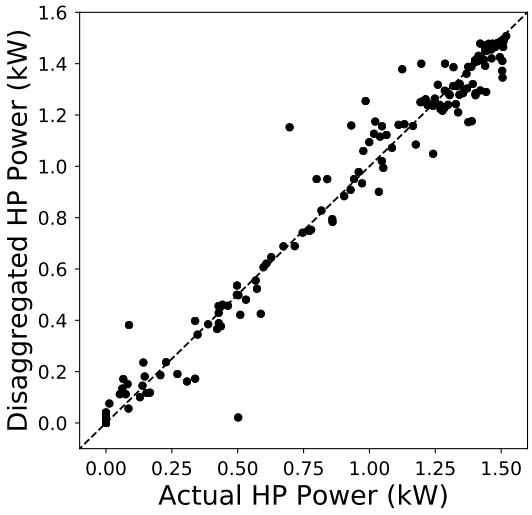


Figure 3: Heat pump disaggregation test performance. Since there is no significant over- or underprediction, very little error is introduced to the heat pump power model training.

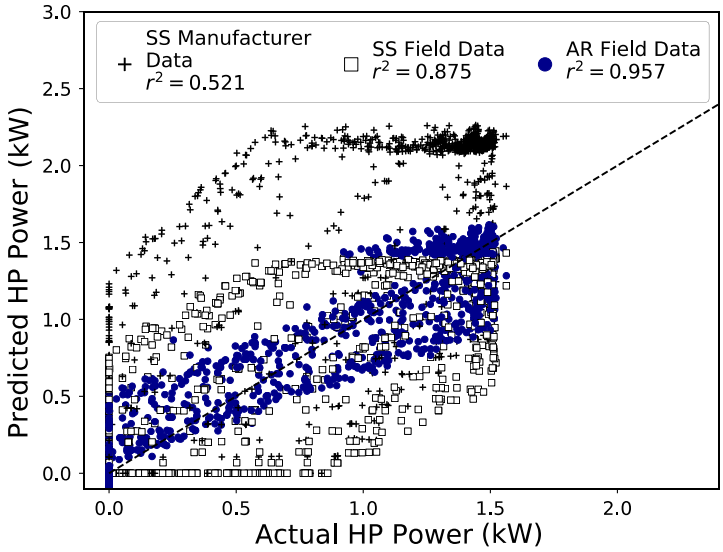


Figure 4: Heat pump power model results. The AR approach accounts for transient effects like control delays and warm-up times and thus predicts significantly better, particularly when the power is not yet at its steady-state. Overprediction from the manufacturer data is likely due to differences in the lab and field installations.

### *2.3. Building Model Identification*

Again, due to the scalability and accessibility of data-driven methods, we use data to determine the building model. While reduced-order white-box building models based on building construction perform well in control [31], they often require detailed simulation models in software like EnergyPlus [32], severely limiting their potential for scalability. In contrast, black-box models contain no explicit knowledge of the underlying thermodynamical equations or building construction, but instead infer the temperature dynamics entirely through data. Linear autoregressive models, which consist of linear combinations of features from previous timesteps, are used in [33] to model both energy consumption and indoor temperature. More complex machine learning techniques such as random forest [34] and neural networks [35] can produce even more accurate models. Common input features to these models include indoor and outdoor temperature, wind speed, solar radiation, occupancy, and internal electrical loads. However, black-box models are at serious risk of overfitting if the training data are not sufficiently diverse.

Buildings often have either periodic or stagnant setpoint schedules which can cause insufficient system excitation leading to overfitting. If data is insufficient and model input features are not carefully engineered, the model could learn the setpoint schedule and closed-loop control dynamics, rather than the temperature's response to control inputs. For example, the model can predict the indoor temperature to oscillate between 20°C and 21°C during daytime, regardless of weather conditions or heat pump input, since this is what occurs throughout the entire training dataset. Moreover, this overfitting is not immediately apparent when used on test data from the same distribution, as the control algorithm and setpoint schedule can remain the same, and the model is not asked to generalize to other potential control methods like MPC.

Most studies remedy this issue through pseudorandom control inputs over a variety of weather conditions. This is easily done in simulation studies such as in [36] and [37], where buildings can be subjected to any condition to generate training data. For real buildings, this method is considerably more difficult due to occupant comfort requirements and uncontrollable weather patterns. Ref. [38] generates sufficient training data by implementing an identification period during the weekend for an experimental unoccupied commercial building in spring and summer. While this process is technically feasible for residential buildings when residents are not home, giving up control of their heating or cooling system for long enough to generate a sufficient amount of training data may not be ideal for many people. A more recently applied method to prevent overfitting is through transfer learning using neural networks [39]. This method uses a different simulated building to generate enough training data to learn the underlying thermodynamics, requiring significantly less data from the building to be modeled. However, neural networks are nonlinear and nonconvex, making them extremely difficult to implement in MPC.

Because of these issues, we opt for a grey-box modeling approach that uses data to fit parameters for a linear reduced-order physics-based model to reduce the amount and diversity of training data needed. For buildings, the physical model almost always takes the form of a thermal circuit involving resistors and capacitors, and can range from single order 1R1C (1 resistor, 1 capacitor) models for small, lightweight buildings [40], to higher-order models with dozens of resistors and capacitors in large multi-zone buildings [41].

Determining optimal parameters is a nonconvex, nonlinear optimization problem making the globally optimal solution difficult to obtain. To circumvent this issue, some studies attempt to find globally optimal solutions using metaheuristic optimization techniques like genetic algorithms [42] and particle swarm optimization [43]. Others use Latin hypercube sampling to generate many initial guesses for convex optimization software to search for global minima [44]. However, finding a globally optimal solution is highly dependent on the chosen bounds and initial guesses. Without *a priori* knowledge of the building construction, guessing tight bounds without excluding the optimal solution can be challenging.

Another difficulty is that data must be collected from sensors that capture enough of the building's state to fit the chosen order of the model. Typical RC modeling studies, such as [45], have extensive knowledge of floor plans, temperatures, and heat gains throughout the building. Installing these sensors can be confusing, costly, and intrusive, though, reducing the scalability potential. For higher-order models, a large number of data sources and system excitation are required to prevent overfitting and subsequent control performance degradation [46] . Therefore, it is important to choose a model that is sufficiently complex to capture the building's dynamics while still considering the limitations based on the amount of data sources and system excitation available.

Finally, when identifying models for use in model predictive control, it is vital to determine the optimal system parameters based on multi-step prediction rather than on next-step prediction and recursively predicting future timesteps. For systems with low state excitation, such as building

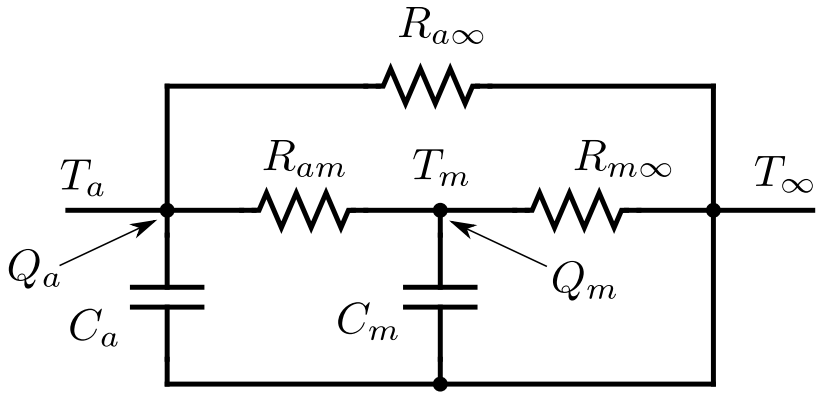


Figure 5: Three Resistor two capacitor (3R2C) thermal circuit representing the grey box building model. The subscript $a$ represents the indoor air and other low thermal capacitance materials, $m$ represents the building's mass, and $\infty$ is the outdoor air.

thermodynamics, next step prediction can produce models that perform unrealistically well on test data. For example, if the temperature does not change much between timesteps, even a persistence model can perform remarkably well when only tested on its next-step prediction. In order to capture the evolution of the thermodynamics for MPC, the multi-step prediction horizon should be on the same order as the MPC horizon.

While there has already been a significant amount of research into various RC model types and the algorithms used to determine optimal thermal parameters [44], the novel contribution of this section is a method that can successfully combat overfitting when using only the limited data available from a thermostat, an electricity meter, and a nearby weather station: (1) indoor temperature, (2) outdoor temperature, (3) disaggregated heat pump power, (4) disaggregated indoor power, and (5) solar irradiation. Instead of using an identification period, we only require the system excitation inherently provided by the occupant's own thermostat setpoint schedule, which may allow the temperature to drift during night or when the occupant is away. We then use an iterative bound tightening algorithm combined with convex optimization software to find an optimal solution consistent with expected model error from sensor uncertainty. Finally, we illustrate the discrepancy between next-step and multi-step prediction and compare the generalization performance of our model to common black-box modeling techniques.

*2.3.1. Model Definition*

We take a similar modeling and identification approach to [47], which uses three two-state 3R2C (three resistors and two capacitors) model shown in Fig. 5. The two-state model models the temperature of the indoor air (including any other low thermal capacitance material) and building mass as separate states since the thermal capacitance of the building's mass is often much higher than that of the indoor air. This model is represented by,

$$\begin{aligned} C_a\dot{T}_a(t) &= \frac{T_\infty - T_a(t)}{R_{a\infty}} + \frac{T_m(t) - T_a(t)}{R_{am}} + Q_a \\ C_m\dot{T}_m(t) &= \frac{T_\infty - T_m(t)}{R_{m\infty}} + \frac{T_a(t) - T_m(t)}{R_{am}} + Q_m. \end{aligned} \tag{2}$$

Here, the subscript $a$ refers to the indoor air, the subscript $m$ refers to the building mass, and $\infty$ refers to the outside air. $R_{ij}$ (°C/kW) refers to the thermal resistance between element $i$ and $j$, and $C_i$ (kJ/°C) is the thermal capacitance of element $i$. The sources of external heat gain for the indoor air node $Q_a$ (kW) involve the heat from the heat pump (which blows hot air into the

indoors), heat from internal gains (which can be estimated from internal power consumption), and any heat coming from solar radiation (which enters through windows). In contrast, the additional heat gains to the building's mass only include the solar irradiation. These gains can therefore be defined as,

$$Q_a = Q_{\mathrm{HP},a} + \eta Q_{e,a} + \alpha_a G, \quad Q_m = \alpha_m G, \tag{3}$$

where the superscript $e$ refers to heat from internal gains and $\eta$ is the fraction of indoor power consumption that results in heat gains. Here we use $G$ (kW/m$^2$) as the solar radiation and $\alpha$ (m$^2$) as the solar absorption factor, or the fraction of solar energy absorbed as heat. While the heat transfer from a heat pump can often be modeled using manufacturer performance data, for similar reasons that we describe in Sec. 2.2, the heat pump does not always perform the same in the field as it does in the lab. Therefore, we assume the heat pump COP can be approximately modeled as linearly dependent on the difference between indoor and outdoor air temperature [48] and seek to find the optimal model parameters $\beta_i$. This model gives the heat output $Q_{\mathrm{HP},a}$ as,

$$Q_{\mathrm{HP},a} = \Big(\beta_1(T_\infty - T_a) + \beta_2\Big) P_{\mathrm{HP}}. \tag{4}$$

where $P_{\mathrm{HP}}$ is the disaggregated power of the heat pump.

In order to formulate a complete model for use in control, it is useful to convert Eqs. 2 through 3 from continuous time into a first-order Euler discrete-time state space model of the form,

$$\boldsymbol{x}_{k+1} = \boldsymbol{A}\boldsymbol{x}_k + \boldsymbol{B}_k\boldsymbol{u}_k + \boldsymbol{C}\boldsymbol{w}_k, \tag{5}$$

where,

$$\begin{aligned}
\boldsymbol{x} &= \begin{bmatrix} T_{a,k} \\ T_{m,k} \end{bmatrix}, \quad \boldsymbol{u_k} = \Big[P_{\mathrm{HP},k}\Big], \quad \boldsymbol{w_k} = \begin{bmatrix} T_{\infty,k} \\ Q_{e,k} \\ G_k \end{bmatrix} \\
\boldsymbol{A} &= \begin{bmatrix} 1 - \frac{\Delta t}{C_a}\Big(\frac{1}{R_{a\infty}} + \frac{1}{R_{am}}\Big) & \frac{\Delta t}{C_a R_{am}} \\ \frac{\Delta t}{C_m R_{am}} & 1 - \frac{\Delta t}{C_m}\Big(\frac{1}{R_{m\infty}} + \frac{1}{R_{am}}\Big) \end{bmatrix}, \\
\boldsymbol{B}_k &= \begin{bmatrix} \frac{\Delta t}{C_a}\Big(\beta_1(T_{\infty,k} - T_{\mathrm{set},k}) + \beta_2\Big) \\ 0 \end{bmatrix}, \\
\boldsymbol{C} &= \begin{bmatrix} \frac{1}{R_{a\infty}C_a} & \frac{\eta\Delta t}{C_a} & \frac{\alpha_a \Delta t}{C_a} \\ \frac{1}{R_{m\infty}C_m} & 0 & \frac{\alpha_m \Delta t}{C_m} \end{bmatrix}.
\end{aligned} \tag{6}$$

Here, $k$ represents the timestep and ranges from 0 to the number of training samples $K$. Since the difference between the predefined setpoint $T_{\mathrm{set},k}$ and $T_{a,k}$ must be very small, we use $T_{\mathrm{set},k}$ as a proxy for $T_{a,k}$ to prevent nonlinearity in the control model that results from multiplication of control and state variables.

In contrast to [47], which identifies the consolidated entries of the state-space matrix, we opt to find the actual thermal parameters of the model. While [47] had sufficient excitation to reduce overfitting, we observed that using our data to find the optimal entries in the state-space matrices performed very similarly to black-box linear regression, which we later show is unable to generalize for control when given low system excitation. Finding the individual thermal parameters instead of the matrix entries provides the model with both interpretability and the ability to add physically intuitive constraints on the thermal parameters, such as non-negativity.

These thermodynamic equations can be combined into a nonlinear, nonconvex constrained optimization problem that minimizes the sum of the squared errors between the actual indoor temperature $T_a^k$ and the predicted indoor temperature $\hat{T}_a^k$.

$$\min_{J} \quad \sum_{k=100}^{K} (\hat{T}_a^k - T_a^k)^2 \tag{7a}$$

subject to

$$\boldsymbol{x}_{k+1} = \boldsymbol{A}\boldsymbol{x}_k + \boldsymbol{B}_k\boldsymbol{u}_k + \boldsymbol{C}\boldsymbol{w}_k \quad \forall k \in K \tag{7b}$$

$$\underline{b}^i \leq J_i \leq \bar{b}^i \quad \forall i \in J \tag{7c}$$

$$\hat{T}_a^0 = \hat{T}_m^0 = T_a^0. \tag{7d}$$

$$\hat{T}_a^{im} = T_a^{im} \quad \forall i \in \{1..K/m\} \tag{7e}$$

In Eq. 7c, we add conservative upper and lower constraints on the variable bounds, denoted as $\underline{b}^i$ and $\bar{b}^i$ and indexed over the set of decision variables $J$, shown in Tab. 1. In Eq. 7d, we assume the initial predicted mass and air temperature are equal to the initial air temperature, but only consider errors from after the first 100 timesteps to eliminate the effects of initial condition assumptions. Eq. 7e imposes the multistep prediction constraint: Knowledge of the true temperature is only given every $m$ timesteps.

Table 1: Optimization parameters used for determining model parameters

| Parameter | Value | | Units |
|---|---|---|---|
| Decision Variable Bounds | Lower | Upper | |
| $R_{am}$, $R_{a\infty}$, $R_{n\infty}$ | 0 | 40 | $^\circ C$/kW |
| $C_a$, $C_m$ | 0 | 1e5 | kW/$^\circ C$ |
| $C_m$ | 0 | 5e5 | kW/$^\circ C$ |
| $\alpha_a$, $\alpha_m$ | 0 | 5 | $m^2$ |
| $\hat{T}_a$, $\hat{T}_m$ | 10 | 35 | $^\circ C$ |
| $\beta_1$ | -.5 | 0 | $1/^\circ C$ |
| $\beta_2$ | 1 | 5 | |
| $\eta$ | 0 | 2 | |
| Prediction Horizon $m$ | 48 | | |

*2.3.2. Iterative Bound Tightening*

While nonconvex optimization solvers such as SCIP [49], Couenne [50], and Baron [51] can use additional heuristic methods and spatial branch and bound to search for global optimal solutions, these solvers are unable to converge to an optimal solution due to the problem size (appr. 30,000 constraints for 5000 data samples). In contrast, the convex optimization solver Ipopt [52] can quickly find a locally optimal solution, but there is no guarantee that it can find the global optimum. In fact, Ipopt consistently converges to a local minimum where either $R_{am}$ or $R_{a\infty}$ are set to their upper bound, somewhat regardless of the initial guess and how high the upper bounds are set. This phenomenon turns out to be very similar to how black-box methods overfit on this data: Since the

indoor temperature remains relatively constant for much of the time, a good solution can be found by putting an unrealistically high weight on the previous indoor temperature (equivalent to very high resistance values in the thermal circuit model).

To prevent this overfitting, we implement an iterative bound tightening (IBT) algorithm that tightens each active upper bound until the model performance on validation data aligns with the expected lower bound on model error, which is derived from thermostat sensor uncertainty. In this way, we borrow ideas from hyperparameter tuning in machine learning algorithms. In essence, iterative bound tightening performs an efficient grid search of the optimal bounds, or hyperparameters, that provide good training performance and prevent overfitting. In addition, since the problem starts with new random initial points over a reduced search space on each iteration, we increase the chance that the solver finds a global minimum. The algorithm is as follows:

**Algorithm 1:** Iterative Bound Tightening (IBT)

```
Solve problem with conservative upper and lower bounds;
Initialize list B of model parameters that are at their bound;
while True do
    for bound in B do
        Tighten bound by some amount β;
        Solve problem with new bounds;
        Calculate validation error;
    end
    if no model parameters at bounds or error is diverging then
        Stop;
    else
        Update B with new model parameters at new bounds;
    end
end
Choose bounds with lowest validation model error;
Retrain final model on validation and training data.
```

Fig. 6 shows the results of the bound tightening algorithm (IBT) on our experimental data. Increasingly higher bounds on $R_{a\infty}$ and $R_{m\infty}$ cause the model to plateau at the same RMSE on training data. However, when tested on validation data, the RMSE consistently increases at higher upper bounds, indicating overfitting. The final parameter bounds were taken from the model that gave the lowest validation data error. Using these optimal bounds, the model was trained on both the validation and training data and then tested on test data to present the final IBT model error.

*2.3.3. Building Model Identification Results*

To show the performance of the IBT multistep prediction model, we compare it to other modeling methods used in literature: linear regression (LR) [53], random forest regression (RF) [34], and artificial neural networks (ANN) [39]. Each of these black-box models use the same input data as the IBT model but also include autoregressive terms. In this comparison, we seek to show the generalizing capability of the IBT model as well as the importance for training the model for multi-step prediction. Tab. 2 shows each of the models' performance over the MPC horizon on test data when trained for next-step prediction and multi-step prediction. Next-step prediction models were therefore used recursively, where the predicted output at each timestep becomes part of the input for the next timestep's prediction.

For each method, training the model to minimize the error over the entire control horizon

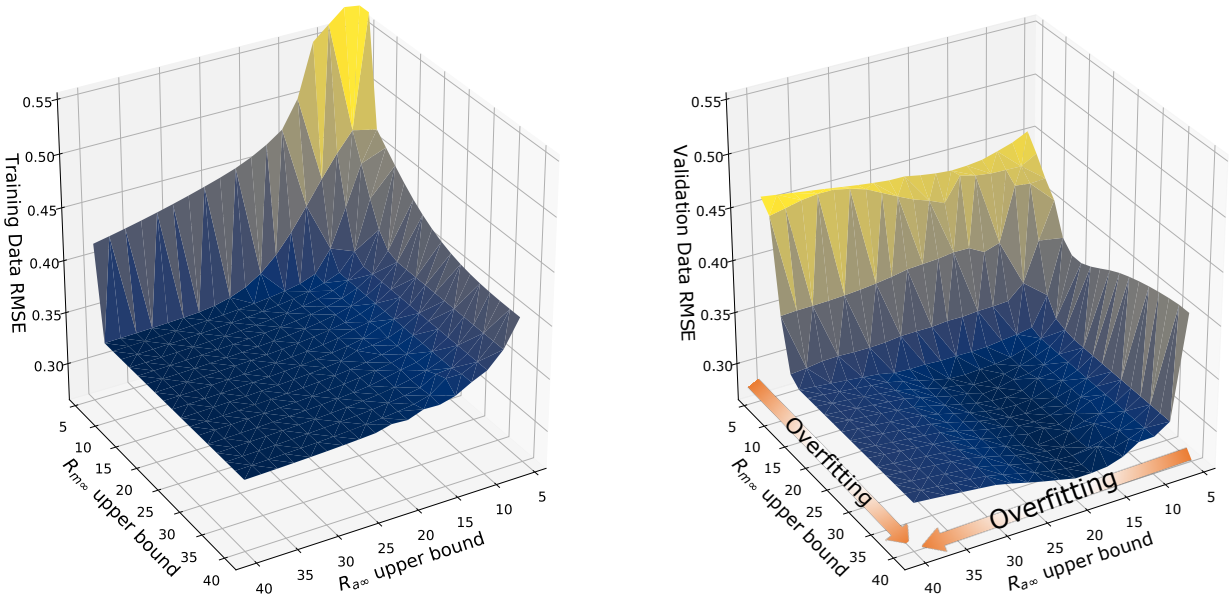


Figure 6: Building model identification RMSE for training data (left) and validation data (right). While higher bounds on $R_{a\infty}$ and $R_{m\infty}$ produce lower training error, it begins to overfit, resulting in error increase on validation data.

Table 2: Building model test RMSE when used to predict over the MPC horizon. Next-step models were used to predict each step recursively, with the prediction for each step as the input for the next step's prediction. While LR provides the lowest RMSE, it has overfit, explained in Fig. 7

| Model | Trained for next-step prediction ($^\circ C$) | Trained for multi-step prediction ($^\circ C$) |
|---|---|---|
| IBT | .468 | .285 |
| LR | .197 | .190 |
| RF | .417 | .349 |
| ANN | .761 | .301 |

introduces significantly lower error into the MPC compared to minimizing the next step error and predicting each next step recursively. However, model error does not tell the entire story. While the LR multistep model provides the lowest error on test data, further analysis of the model coefficients shows that this black-box method disproportionately weights the previous timestep's temperature and gives very little weight to control inputs. This overfitting phenomenon arises in every black-box model we tested and is illustrated in Fig. 7 through the forced and unforced responses for each model on a winter day. Despite having significantly lower test-error, the LR multi-step model cannot generalize to new control schemes and predicts the indoor temperature to gradually rise in response to zero heat input. In contrast, the IBT multistep model, with its underlying thermodynamic equations, can provide low test error while at the same time be used in model-based control.

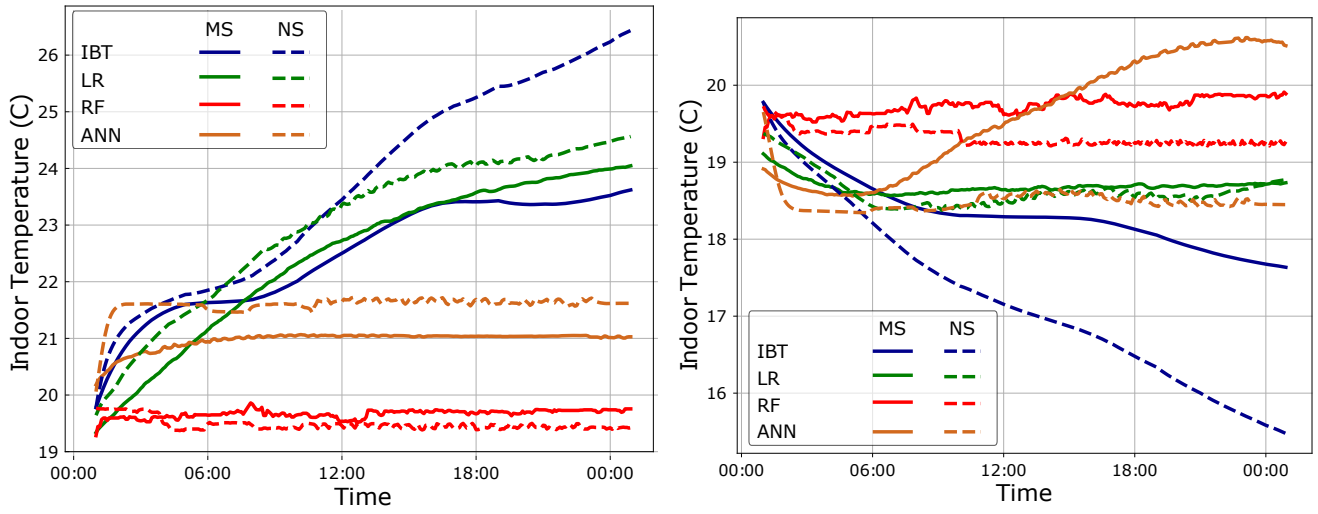


(a) (b)

Figure 7: The forced (a) and unforced (b) responses over a winter day for each control method trained for both multi-step (MS) and next-step (NS) prediction. While IBT can generalize to new control inputs and new indoor temperature ranges, the black-box methods (LR, RF, ANN) do not learn the underlying physics and cannot be used in MPC.

### *2.4. Disturbance Prediction*

#### *2.4.1. Weather Prediction*

Future weather conditions over the control horizon can be reliably obtained from dedicated forecasting services to a high degree of accuracy. To emulate these forecasts for use in our simulation, we use data-driven models trained on historical weather data. The two weather related disturbances we consider in our model are outdoor air temperature and solar insolation, since these have the largest effect on indoor heating or cooling load.

Solar insolation and outdoor air temperature are predicted using linear autoregressive models based on input features readily available from a nearby weather station. These include solar insolation, dewpoint, wind speed, pressure, and time of day. The output of the models are the next four hours of outdoor temperature and solar insolation. Training data comes from [21] for the winters (Dec. to Mar.) of 2016 and 2017 in Ithaca, totalling $14,000$ samples. The models are then tested on the winter of 2018. RMSE for outdoor temperature prediction is 0.99 $^{\circ}C$, while RMSE for solar radiation is 0.09 kW/m$^2$.

#### *2.4.2. Aggregate Load Prediction*

The peak load seen by the grid is the aggregate whole-home power, which includes both the centrally controllable heat pump power and the internal loads each resident uses. Therefore, in order for the controller to adequately plan the heat pump control to minimize the aggregate peak power, the internal load profile must also be forecasted. While predicting single building internal loads proves quite difficult due to its volatility [54], the aggregate internal load profile of many homes contains identifiable trends and patterns and allows an aggregate controller to have improved disturbance predictions.

Training data for aggregate load prediction comes from our measurements of non-heat pump related electricity consumption in three different residences similar to the building described in Sec. 2. A load profile representing the power consumption of more than just three units is created by augmenting this data. For each additional unit added to the load profile, we randomly shift

each individual load profile by up to three hours and up to 15 days to simulate different resident behaviors. Training and test data come from different combinations of these random shifts and contain 6000 samples each.

Again, we use a linear autoregressive model for forecasting. The previous 6 hours, or 72 timesteps, of the power profile as well as the time of day were used as inputs to the model, with the next four hours, or 48 timesteps as the output. We find that while more complex machine learning methods like random forest regression and neural networks provide lower RMSE, they tend to underpredict high peak loads because these loads happen very rarely in the training dataset. If high peak loads are not accurately predicted, peak load minimization MPC does not have sufficient information to plan an optimal control sequence and will result in higher peak load costs. Therefore, we use linear regression because it predicts these peaks much better, particularly if they occur in the next hour. This results in an RMSE of 3.01 kW, or around 15% of the aggregate peak of 24 residences.

## 3. Control Problem Formulation

We now implement these data-driven models into MPC. MPC incorporates a model of the system and predictions of future disturbances into an optimization problem that minimizes the control objective over the desired control horizon. At each timestep, only the optimal control for the first timestep is implemented. At the next timestep, this process is repeated with the updated state and disturbance predictions. Different cost functions and constraints can be added to the formulation, allowing the controller to be applicable to many types of heat pump systems and consider a wide variety of control objectives.

### *3.1. Constraints*

#### *3.1.1. Thermal comfort*

We define thermal comfort as a temperature range $T_\delta$ above and below the thermostat setpoint $T_{\text{set}}$. Since setpoints are assumed to be customizable by the resident, violations outside of this thermal comfort band are penalized to maintain feasibility. Residents often use setpoint setbacks in the interest of reducing energy consumption, thus causing large unavoidable comfort violations. Since these violations are expected and in the primary interest of energy efficiency rather than actual comfort, we use different values for upper and lower comfort violations, denoted as $\overline{\pi}_{\text{pen}}$ and $\underline{\pi}_{\text{pen}}$ (\$/s/°C), respectively. These violations are enforced by the constraints,

$$\begin{aligned} T_j^i &\leq T_{\text{set},j}^i + T_{\delta,j}^i + \overline{T}_{\text{pen},j}^i \quad \forall j \in N, i \in B \\ T_j^i &\geq T_{\text{set},j}^i - T_{\delta,j}^i - \underline{T}_{\text{pen},j}^i \quad \forall j \in N, i \in B. \end{aligned} \tag{8}$$

Here, $T_j^i$ is the indoor temperature of the $i$th building at timestep $j$, $\underline{T}_{\text{pen},j}^i$ is the lower comfort violation, and $T_\delta^i$ is the occupant's comfort band above or below the setpoint. $N$ represents the prediction horizon, and $B$ is the set of buildings. Note that the comfort band can vary based on time of day as well. For the purpose of this simulation, we assume $T_{\delta,j}^i$ is 4 °C during work hours (8:00 am to 4:00 pm), 2 °C at night (11:00 pm to 6:00 am), and 1 °C all other times.

#### *3.1.2. Minimum on and off times*

To prevent compressor short-cycling and the resulting performance degradation, we enforce minimum on and off times for each heat pump. This means that once a heat pump turns on (or off) it must remain on (or off) for a minimum amount of time. To do this, we introduce the binary variables $v_j^{\uparrow,i}$ and $v_j^{\downarrow,i}$ to indicate whether the heat pump in the $i$th building turned on or

off, respectively, at timestep $j$. The values of these variables and the minimum on and off times are enforced by the following constraints,

$$u_j^i - u_{j-1}^i = v_j^{\uparrow,i} - v_j^{\downarrow,i} \quad \forall j \in N, i \in B \tag{9}$$

$$\sum_{k=j-t_{\text{min on}}}^{j} v_k^{\uparrow,i} \leq u_j^i \quad \forall j \in N, i \in B \tag{10}$$

$$\sum_{k=j-t_{\text{min off}}}^{j} v_k^{\downarrow,i} \leq 1 - u_j^i \quad \forall j \in N, i \in B. \tag{11}$$

Here, the minimum on time is denoted by $t_{\text{min on}}$ and the minimum off time is denoted by $t_{\text{min off}}$, which we set equal to 3 (15 minutes).

*3.1.3. Peak Load*

In order to minimize the aggregate peak load, we introduce another decision variable $P_{\max}$ through the constraint,

$$P_{\max} \geq Q_{\text{e},j}^{\text{agg}} + \sum_{i \in B} P_{\text{HP},j}^i \quad \forall j \in N. \tag{12}$$

Here, $P_{\max}$ must be greater than or equal to the sum of the heat pump power over all buildings and the aggregate indoor load $Q_{\text{e},j}^{\text{agg}}$ predicted by the model introduced in Sec. 2.4.2.

*3.2. Control objectives*

While most heat pump MPC for improved efficiency and costs is on a local level, the scalability of this approach supports aggregate control to provide local benefits as well as provide some aggregate service to the grid, such as peak load management. In order to show this value of aggregation and the benefit of aggregate MPC, we analyze three potential controllers: aggregate MPC for peak load reduction, local MPC using time-of-use rates, and standard rule-based thermostat control using programmable thermostat setbacks.

*3.2.1. Aggregate Peak Load MPC*

This control policy minimizes the aggregate electrical cost of the buildings using an aggregate peak demand charge pricing structure. This structure charges a delivery charge (\$/kWh) for the amount of energy used as well as a demand charge (\$/kW) based on the highest 5-minute averaged aggregate demand for all buildings. While a demand charge structure is usually only offered to commercial buildings due to their size and controllability, this aggregate control can expand this structure's benefits of discouraging peak load and reduced capacity requirements to an aggregation of residences as well.

This price structure is included in the objective function as,

$$\min_{u_j^i} \quad C = \sum_{j=0}^{N-1} \sum_{i \in B} \left[ \pi_{e,j} P_j^i + \underline{\pi}_{\text{pen}} \underline{T}_{\text{pen},j}^i + \overline{\pi}_{\text{pen}} \overline{T}_{\text{pen},j}^i \right] \Delta t + r \pi_{\max} P_{\max}. \tag{13}$$

Here, $\pi_{e,j}$ is the time-varying delivery charge and $\pi_{\max}$ is the demand charge. In order to properly scale the demand charge for the MPC horizon, we include a scaling parameter $r$ equal to the ratio of the number of minutes in the time horizon to the number of minutes in the month.

*3.2.2. Local Time-of-use MPC*

This control policy uses a residential TOU rate in the objective function and optimizes each building separately, with no communication or coordination between buildings. Utilities use TOU rates to discourage homes from contributing to the aggregate peak load by charging higher prices for on-peak hours. For this simulation, we use New York State Electric and Gas's (NYSEG) service class 12 winter TOU rate [55], where on-peak lasts from 7:00 am to 10:00 am and from 5:00 pm to 10:00 pm, mid-peak lasts from 10:00 am until 5:00 pm, and from 10:00 pm until 11:30 pm, and off-peak is all other times.

The objective function for this rate structure is,

$$\min_{u_j} \quad C = \sum_{j=0}^{N-1} \left[\pi_{e,j} P_j + \underline{\pi}_{\text{pen}} \underline{T}_{\text{pen},j} + \overline{\pi}_{\text{pen}} \overline{T}_{\text{pen},j}\right] \Delta t. \tag{14}$$

*3.2.3. Rule-based hysteresis control with programmable thermostat setbacks*

This control policy represents the most common rule-based control method for smart thermostats. In heating mode, when the indoor temperature falls below the thermostat's lower comfort level, the heat pump will turn on to heat the air until it reaches the upper part of the deadband. At that point, it will turn off and let the temperature drift down again, repeating the cycle. By using custom setpoint schedules, this method represents an occupant's best attempt at balancing thermal comfort with energy efficiency. In addition, energy saving tactics employed by the smart thermostat, such as setpoint setbacks when the user is away, are also reflected in the setpoint schedules used in the simulation. This process can be represented by,

$$u^i_{j+1} = \begin{cases} m, & T^i_j < T^i_{\text{set},j} - T^i_\delta \\ 1-m, & T^i_j > T^i_{\text{set},j} + T^i_\delta \\ u^i_j, & \text{otherwise} \end{cases}. \tag{15}$$

where $m$ is equal to zero if in cooling mode and one if in heating mode.

## 4. Simulation

*4.1. Configuration*

We simulated five days of heating for 24 randomly generated residences and heat pumps, where model parameters were randomly chosen from a uniform distribution of ±20% around the parameters derived from the data in Sec. 2. To introduce model uncertainty, we add variation between the plant and control model parameters such that the average error between the MPC and the plant models is equal to the test error results from model identification. Weather data range from Feb. 24 to Feb. 28 for Ithaca, NY. Both the weather and aggregate load data are subsets of the test data used for evaluating the performance of the prediction models. Twenty-four setpoint schedules were randomly chosen from the collection of over 2000 participants in NY state in the Ecobee Donate your Data dataset [19].

The MPC uses a five-minute timestep with a optimization horizon of 48 steps, or 4 hours. Various optimization and simulation parameters are shown in Tab. 3. We carry out the optimization using the CPLEX 12.9 solver in the Pyomo modeling environment [56] with a maximum run time of five minutes, equal to the timestep length. Since conditions change relatively little over a single timestep, we use each optimal solution as a warm-start for the next timestep to aid solver performance.

Table 3: Parameters used for MPC optimization and numerical simulation

| Parameter | Symbol | Value | Ref. |
|---|---|---|---|
| **Peak Demand Pricing Rate** | | | [55] |
| Demand Charge | $\pi_{\max}$ | \$9.43 /kW | |
| Delivery Charge | $\pi_{e,j}$ | \$0.025 /kWh | |
| **TOU Pricing** | $\pi_{e,j}$ | | [55] |
| On-peak | | \$0.0552 /kWh (7:00-10:00, 17:00-22:00) | |
| Mid-peak | | \$0.0468 /kWh (10:00-17:00, 22:00-23:30) | |
| Off-peak | | \$0.0376 /kWh (23:30-7:00) | |
| **Optimization Parameters** | | | |
| MPC Horizon | $N$ | 4 h | |
| Time Step | $\Delta t$ | 5 min | |
| Number of Buildings | $B$ | 24 | |
| Min HP On time | $t_{\text{min on}}$ | 15 min | |
| Min HP Off time | $t_{\text{min off}}$ | 15 min | |
| Upper comfort violation | $\overline{\pi}_{\text{pen}}$ | \$0.6/°C/h | |
| Lower comfort violation | $\underline{\pi}_{\text{pen}}$ | \$1.2/°C/h | |

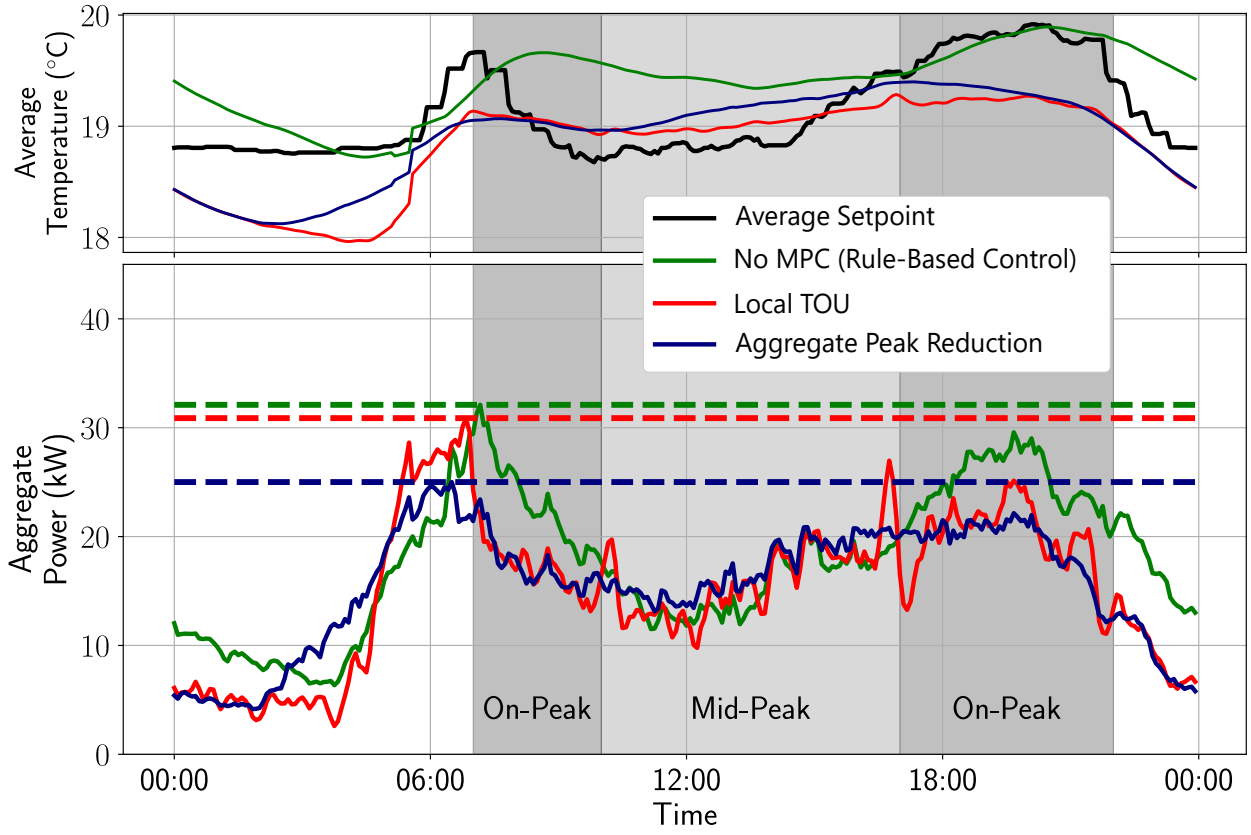


Figure 8: Average daily profiles for indoor temperature, setpoint, and aggregate power for the 5-day simulation in February for each control method. The dotted line shows the average daily peak power for both the heat pump and indoor loads.

*4.2. Controller Performance*

Fig. 8 shows the average daily profiles for indoor temperature, setpoint, and aggregate power for each control algorithm. The dotted line shows their average daily peak power. In aggregate, real user thermostat setpoint schedules play an enormous role in creating peak loads and illustrate the potential unintended consequences of widespread use of smart thermostats. The mean thermostat setpoint, shown by the black line in the upper plot, creates two peaks at around 7:00 am and 7:00 pm, coinciding with on-peak hours. In this case, the programmable thermostat setbacks enabled by smart thermostats can actually exacerbate peak loads at exactly the times the grid would like to discourage increases in demand. During winter in northern climates these peaks occur just before the sun rises and after the sun sets and therefore cannot directly use solar energy, further underscoring the need to reduce peak loads during those times.

While local TOU-based MPC provides a marked reduction in electrical consumption and costs shown in Fig. 9a, it is still unable to reduce the aggregate peak load. Instead of peaking during on-peak hours, the aggregate power profile creates a new peak just before the on-peak hours occur, with a significant drop-off just after. While this may contribute toward flattening the peak if only a fraction of homes used TOU-based control, should this control method become widely available, TOU rates could have significant adverse effects just before on-peak hours begin.

The aggregate demand charge-based MPC is able to consider both other heat pump operation as well as the aggregate non-heat pump electrical demand. The result is a 19% reduction in average daily peak demand and a significantly more even demand profile when compared to TOU-based MPC. Moreover, the aggregate controller can significantly reduce peak demand *without* a significant increase in electrical consumption, shown in Fig. 9a. When compared to the local MPC, the cost of a 19% reduction in average daily peak demand is only a 4.6% increase in energy consumption, primarily during the night when non-heat pump demand is very low and the wholesale electricity price is often cheap.

Costs for each controller under each rate structure are shown in Fig. 9a and are normalized by the cost of the baseline rule-based controller. The aggregate controller has consistently good performance under each rate structure, and provides a consistent savings of 23% when compared to the custom setpoint thermostat control. The local controller provides slightly better savings of 26% on the consumption-based rate structures, but only a 9% savings on the demand charge structure.

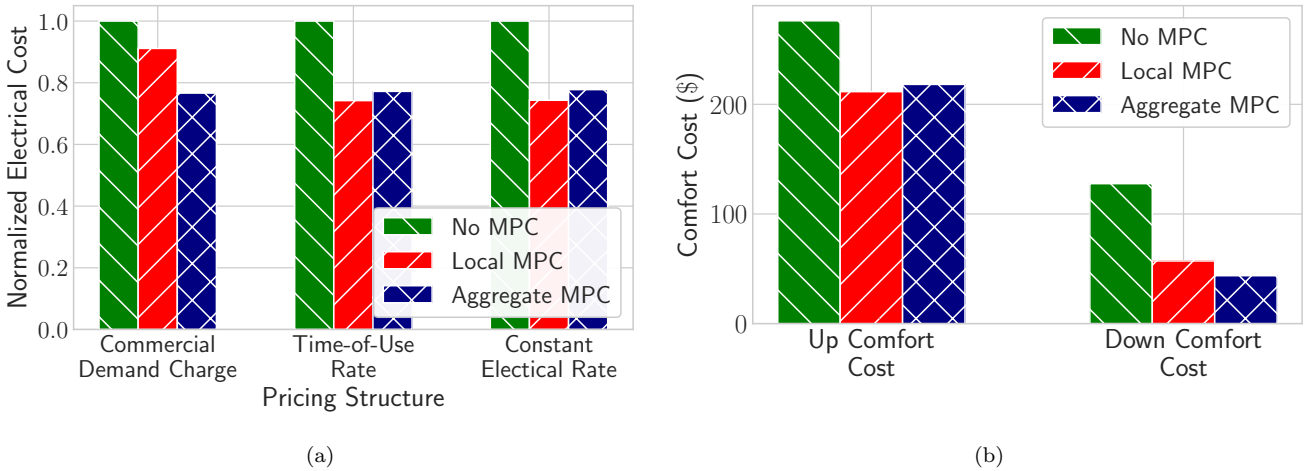


Figure 9: Total control performance results. In (a), the aggregate MPC can provide equivalent cost savings to local MPC under a TOU rate, with the added effect of reducing aggregate peak load. In (b), the aggregate MPC can accomplish this without significant change in thermal comfort.

*4.3. Thermal Comfort*

Fig. 9b shows each algorithm's thermal comfort constraint violation for being above (up) or below (down) the user-defined comfort level. This cost is the portion of the objective function related to thermal comfort constraint violation and is calculated as the product of the thermal comfort penalty $\pi_{\text{pen}}$ and the amount the indoor temperature exceeds the comfort band $T_{\text{pen}}$. Both MPC approaches provide significantly lower thermal comfort violations than rule-based control because of the ability to pre-heat or pre-cool the building so that the temperature is in the correct range prior to setpoint changes. Up comfort violations are much higher because the heat pump can only heat during winter, and thus has no way to lower the temperature in response to large setpoint reductions. However, since the most likely reason for these user-defined thermostat setbacks is in an attempt to save energy and money, up comfort violations are less important than down comfort violations in winter. Since there is no significant difference in thermal comfort between the local and aggregate approaches, we see that the aggregate controller can effectively reduce peak load without any reductions in thermal comfort. Furthermore, the aggregate controller does not unfairly penalize some users to reduce the overall objective, with each unit's aggregate control thermal comfort violation changing less than 5% compared to the local control.

## 5. Conclusion

In this paper, we presented a plug-and-play methodology for using only two smart-home devices – a smart thermostat and smart electricity meter – to provide scalable system identification and model predictive control (MPC) for heat pump and smart grid interactivity. This methodology uses the devices already present in a growing number of homes to model the building and heat pump, as well as provide optimal control for improved comfort and grid services. Since it requires no additional hardware implementation or visit from a professional, this method is scalable and non-intrusive, considerably reducing barriers to adoption.

While we have introduced two new modeling methods that make data-driven building and heat pump modeling feasible when limited data sources are available, there are inevitably some cases where the available data is still insufficient. Examples include thermostats with little setpoint variation or large homes where one thermostat temperature can not accurately represent the whole home. In these cases, adaptive control techniques that artificially excite the system or multiple thermostats and changes to the building model could be necessary. Moreover, the effects of different indoor loads on the indoor temperature can vary widely and the associated parameter $\eta$ is often difficult to estimate. For example, large loads like electric vehicles may need to be submetered or separated using energy disaggregation to distinguish them from heat producing loads inside the home.

Our results show that while smart thermostat setpoint schedules and local MPC can reduce energy consumption, they can also have unintended consequences on the aggregate energy system by creating new or more pronounced peaks in demand during or just before on-peak hours. The scalability of our approach enables aggregate MPC to mitigate these aggregate effects on the energy system, while still providing local benefits of improved thermal comfort and reduced energy costs. Results show that this aggregate MPC can reduce the average daily peak load by 19% without a significant increase in energy consumption or decrease in thermal comfort when compared to a local approach. Therefore, with this scalable approach, we can provide an easily implemented and improved control for smart thermostats and heat pumps, while at the same mitigating the effects that an aggregation of units may have on the overall energy system.

**Acknowledgments**

The authors acknowledge the support from the National Science Foundation (NSF) under grant 1711546, the NSF Graduate Research Fellowships Program (to ZEL) and the Cornell Atkinson Center for Sustainability.

**References**

[1] US Energy Information Administration, Residential Energy Consumption Survey, `https://www.eia.gov/consumption/residential/data/2015/index.php/`, 2018.

[2] K. J. Chua, S. K. Chou, W. M. Yang, Advances in heat pump systems: A review, Applied Energy 87 (2010) 3611–3624.

[3] R. Aldrich, J. Grab, D. Lis, Northeast/Mid-Atlantic air-source heat pump market strategies report 2016 update, Northeast Energy Efficiency Partnerships, 2017.

[4] New York City's Roadmap to 80 x 50, New York City Mayor's Office of Sustainability, 2017.

[5] APPRISE Incorporated, NYSERDA Low- to moderate-income market characterization report, New York State Energy Research and Development Authority (NYSERDA), 2017.

[6] Z. E. Lee, Q. Sun, Z. Ma, J. Wang, J. S. MacDonald, K. Max Zhang, Providing grid services with heat pumps: A review, ASME Journal of Engineering for Sustainable Buildings and Cities 1 (2020).

[7] X. Li, J. Wen, Review of building energy modeling for control and operation, Renewable and Sustainable Energy Reviews 37 (2014) 517–537.

[8] Z. E. Lee, K. Gupta, K. J. Kircher, K. M. Zhang, Mixed-integer model predictive control of variable-speed heat pumps, Energy and Buildings 198 (2019) 75–83.

[9] Y. Zhao, Y. Lu, C. Yan, S. Wang, MPC-based optimal scheduling of grid-connected low energy buildings with thermal energy storages, Energy and Buildings 86 (2015) 415 – 426.

[10] B. Baeten, T. Confrey, S. Pecceu, F. Rogiers, L. Helsen, A validated model for mixing and buoyancy in stratified hot water storage tanks for use in building energy simulations, Applied Energy 172 (2016) 217 – 229.

[11] L. Zhang, N. Good, P. Mancarella, Building-to-grid flexibility: Modelling and assessment metrics for residential demand response from heat pump aggregations, Applied Energy 233-234 (2019) 709 – 723.

[12] J. L. Mathieu, S. Koch, D. S. Callaway, State estimation and control of electric loads to manage real-time energy imbalance, IEEE Transactions on Power Systems 28 (2013) 430–440.

[13] D. S. Callaway, Tapping the energy storage potential in electric loads to deliver load following and regulation, with application to wind energy, Energy Conversion and Management 50 (2009) 1389–1400.

[14] H. Hao, B. M. Sanandaji, K. Poolla, T. L. Vincent, Aggregate flexibility of thermostatically controlled loads, IEEE Transactions on Power Systems 30 (2015) 189–198.

[15] J. Clauß, L. Georges, Model complexity of heat pump systems to investigate the building energy flexibility and guidelines for model implementation, Applied Energy 255 (2019) 113847.

[16] Smart Thermostats Gain Traction in Europe and North America., Berg Insight, 2017.

[17] K. Ehrhardt-Martinez, K. A. Donnelly, J. A. Laitner, Advanced Metering Initiatives and Residential Feedback Programs: A Meta-Review for Household Electricity-Saving Opportunities, American Council for an Energy-Efficient Economy (ACEEE), 2010.

[18] A. Meier, T. Ueno, M. Pritoni, L. Rainer, A. Daken, D. Baldewicz, What can connected thermostats tell us about american heating and cooling habits?, in: ECEEE 2019 Summer Study Proceedings, 2019.

[19] Ecobee Inc., Donate your data, 2019. `https://www.ecobee.com/donateyourdata/`.

[20] N. Mahdavi, J. H. Braslavsky, C. Perfumo, Mapping the effect of ambient temperature on the power demand of populations of air conditioners, IEEE Transactions on Smart Grid 9 (2018) 1540–1550.

[21] M. Sengupta, Y. Xie, A. Lopez, A. Habte, G. Maclaurin, J. Shelby, The National Solar Radiation Data Base (NSRDB), Renewable and Sustainable Energy Reviews 89 (2018) 51 – 60.

[22] C. Verhelst, F. Logist, J. V. Impe, L. Helsen, Study of the optimal control problem formulation for modulating air-to-water heat pumps connected to a residential floor heating system, Energy and Buildings 45 (2012) 43–53.

[23] D. Korn, J. Walczyk, A. Jackson, A. Machado, J. Kongoletos, E. Pfann, Ductless Mini-Split Heat Pump Impact Evaluation, Cadmus, 2016.

[24] K. D. Anderson, M. E. Bergés, A. Ocneanu, D. Benitez, J. M. F. Moura, Event detection for non intrusive load monitoring, in: IECON 2012 - 38th Annual Conference on IEEE Industrial Electronics Society, 2012, pp. 3312–3317.

[25] J. Z. Kolter, T. Jaakkola, Approximate inference in additive factorial HMMs with application to energy disaggregation, in: N. D. Lawrence, M. Girolami (Eds.), Proceedings of the Fifteenth International Conference on Artificial Intelligence and Statistics, volume 22 of *Proceedings of Machine Learning Research*, PMLR, La Palma, Canary Islands, 2012, pp. 1472–1482. URL: `http://proceedings.mlr.press/v22/zico12.html`.

[26] J. Z. Kolter, S. Batra, A. Y. Ng, Energy disaggregation via discriminative sparse coding, in: Proceedings of the 23rd International Conference on Neural Information Processing Systems - Volume 1, NIPS'10, Curran Associates Inc., Red Hook, NY, USA, 2010, p. 1153–1161.

[27] S. Drenker, A. Kader, Nonintrusive monitoring of electric loads, IEEE Computer Applications in Power 12 (1999) 47–51.

[28] A. Rahimpour, H. Qi, D. Fugate, T. Kuruganti, Non-intrusive energy disaggregation using non-negative matrix factorization with sum-to-k constraint, IEEE Transactions on Power Systems 32 (2017) 4430–4441.

[29] F. Pedregosa, G. Varoquaux, A. Gramfort, V. Michel, B. Thirion, O. Grisel, et al., Scikit-learn: Machine learning in Python, Journal of Machine Learning Research 12 (2011) 2825–2830.

[30] Coleman Heating and Air Conditioning, Technical guide lx series split system heat pumps, https://bit.ly/35m8MQd, 2019.

[31] W. J. Cole, E. T. Hale, T. F. Edgar, Building energy model reduction for model predictive control using OpenStudio, in: 2013 American Control Conference, 2013, pp. 449–454.

[32] D. B. Crawley, C. O. Pedersen, L. K. Lawrie, F. C. Winkelmann, Energyplus: Energy simulation program, ASHRAE Journal 42 (2000) 49–56.

[33] J. Ma, J. Qin, T. Salsbury, P. Xu, Demand reduction in building energy systems based on economic model predictive control, Chemical Engineering Science 67 (2012) 92 – 100. Dynamics, Control and Optimization of Energy Systems.

[34] M. W. Ahmad, M. Mourshed, Y. Rezgui, Trees vs Neurons: Comparison between random forest and ANN for high-resolution prediction of building energy consumption, Energy and Buildings 147 (2017) 77–89.

[35] Y. Chen, Y. Shi, B. Zhang, Optimal control via neural networks: A convex approach, 2018. arXiv:1805.11835.

[36] D. Sturzenegger, D. Gyalistras, M. Morari, R. S. Smith, Semi-automated modular modeling of buildings for model predictive control, in: Proceedings of the Fourth ACM Workshop on Embedded Sensing Systems for Energy-Efficiency in Buildings, BuildSys '12, Association for Computing Machinery, New York, NY, USA, 2012, p. 99–106. URL: https://doi.org/10.1145/2422531.2422550.

[37] B. Sun, P. B. Luh, Q. Jia, Z. Jiang, F. Wang, C. Song, Building energy management: Integrated control of active and passive heating, cooling, lighting, shading, and ventilation systems, IEEE Transactions on Automation Science and Engineering 10 (2013) 588–602.

[38] Q. Hu, F. Oldewurtel, M. Balandat, E. Vrettos, D. Zhou, C. J. Tomlin, Building model identification during regular operation - empirical results and challenges, in: 2016 American Control Conference (ACC), 2016, pp. 605–610.

[39] Y. Chen, Z. Tong, Y. Zheng, H. Samuelson, L. Norford, Transfer learning with deep neural networks for model predictive control of HVAC and natural ventilation in smart buildings, Journal of Cleaner Production 254 (2020) 119866.

[40] S. Bashash, H. K. Fathy, Modeling and control insights into demand-side energy management through setpoint control of thermostatic loads, in: Proceedings of the 2011 American Control Conference, 2011, pp. 4546–4553.

[41] J. E. Braun, N. Chaturvedi, An inverse gray-box model for transient building load prediction, HVAC&R Research 8 (2002) 73–99.

[42] G. Platt, J. Li, R. Li, G. Poulton, G. James, J. Wall, Adaptive HVAC zone modeling for sustainable buildings, Energy and Buildings 42 (2010) 412 – 421.

[43] C. Andrade-Cabrera, W. J. N. Turner, D. J. Burke, O. Neu, D. Finn, Lumped parameter building model calibration using particle swarm optimization, 2016.

[44] R. De Coninck, F. Magnusson, J. Åkesson, L. Helsen, Toolbox for development and validation of grey-box building models for forecasting and control, Journal of Building Performance Simulation 9 (2015).

[45] J. Date, J. A. Candanedo, A. K. Athienitis, Control-oriented modelling of thermal zones in a house: a multi-level approach, in: International High Performance Buildings Conference, 2016. URL: `http://docs.lib.purdue.edu/ihpbc/229`.

[46] D. H. Blum, K. Arendt, L. Rivalin, M. A. Piette, M. Wetter, C. T. Veje, Practical factors of envelope model setup and their effects on the performance of model predictive control for building heating, ventilating, and air conditioning systems, Applied Energy 236 (2019) 410–425.

[47] E. Vrettos, E. C. Kara, J. MacDonald, G. Andersson, D. S. Callaway, Experimental demonstration of frequency regulation by commercial buildings Part I: Modeling and hierarchical control design, IEEE Transactions on Smart Grid 9 (2018) 3213–3223.

[48] O. Ruhnau, L. Hirth, A. Praktiknjo, Time series of heat demand and heat pump efficiency for energy system modeling, Scientific Data 6 (2019) 189.

[49] A. Gleixner, M. Bastubbe, L. Eifler, T. Gally, G. Gamrath, et al., The SCIP Optimization Suite 6.0, Zuse Institute Berlin, 2018.

[50] P. Belotti, C. Kirches, S. Leyffer, J. Linderoth, J. Luedtke, A. Mahajan, Mixed-integer nonlinear optimization, Acta Numerica 22 (2013) 1–131.

[51] M. Tawarmalani, N. V. Sahinidis, A polyhedral branch-and-cut approach to global optimization, Mathematical Programming 103 (2005) 225–249.

[52] A. Wächter, L. T. Biegler, On the implementation of an interior-point filter line-search algorithm for large-scale nonlinear programming, Mathematical Programming 106 (2006) 25–57.

[53] P.-D. Moroşan, R. Bourdais, D. Dumur, J. Buisson, Building temperature regulation using a distributed model predictive control, Energy and Buildings 42 (2010) 1445 – 1452.

[54] N. Somu, G. R. M. R, K. Ramamritham, A hybrid model for building energy consumption forecasting using long short term memory networks, Applied Energy 261 (2020) 114131.

[55] New York State Electric and Gas Corporation, Electricity service rate, 2018. Service class 12, Rate No. 115-12-00.

[56] W. E. Hart, C. D. Laird, J.-P. Watson, D. L. Woodruff, G. A. Hackebeil, B. L. Nicholson, et al., Pyomo–optimization modeling in python, volume 67, second ed., Springer Science & Business Media, 2017.