

Applying Extended Kalman Filters to Adaptive Thermal Modelling in Homes

Muddasser Alam, Alex Rogers[†], James Scott^{*}, Kamran Ali⁺ and Frederik Auffenberg[◇]

Department of Engineering Science, University of Oxford; [†]Department of Computer Science, University of Oxford; ^{*}Microsoft Research Cambridge, Cambridge; ⁺Usman Institute of Technology, Pakistan; [◇]Electronics and Computer Science, University of Southampton.
moody@robots.ox.ac.uk; [†]alex.rogers@cs.ox.ac.uk; ^{*}jws@microsoft.com; ⁺kamran.ali@uit.edu;
[◇]fa1c12@ecs.soton.ac.uk

ARTICLE HISTORY

Compiled March 13, 2017

ABSTRACT

Space-heating accounts for more than 40% of residential energy consumption in some countries (e.g., the UK and the US) and thus is a key area to address for energy efficiency improvement. To do so, *intelligent* domestic heating systems (IDHS) equipped with sensors and technologies that minimise user-input, have been proposed for optimal heating control in homes. However, a key challenge for IDHS is to obtain sufficient knowledge of the thermal dynamics of the home to build a thermal model that can reliably predict the spatial and temporal effects of its actions (e.g., turning the heating on or off or use of multiple heaters). This challenge of learning a thermal model has been studied extensively for decades in large purpose-built buildings (such as offices, educational, commercial or communal residential buildings) where machine learning is used to infer suitable thermal models. However, we believe that the technological gap between homes and buildings is fast vanishing with the advent of home automation and cloud computing, and the techniques and lessons learned in purpose-built buildings are increasingly applicable to homes too; with necessary modifications to tackle the challenges unique to homes (e.g., impact of household activities, diverse heating systems, more lenient occupancy schedule). Following this philosophy, we present a methodical study where stochastic grey-box modelling is used to develop thermal models and an extended Kalman filter (EKF) is used for parameter estimation. To demonstrate the applicability in homes, we present the case-study of a room in a family house equipped with underfloor heating and custom-built .NET Gadgeteer hardware. We built grey-box models and use the EKF to infer the thermal model of the room. In doing so, we use our in-house collected data to show that, in this instance, our thermal model predicts the indoor air temperature where the 95th percentile of the absolute prediction error is 0.95°C and 1.37°C for 2 and 4 hours predictions, respectively; in contrast to the corresponding 2.09°C and 3.11°C errors of the existing (historical-average based) thermal model.

KEYWORDS

Heating Systems; Thermal Modelling, Extended Kalman Filter, Underfloor Heating.

1. Introduction

Efficient space-heating systems are an integral part of the transition to a low-carbon economy for governments around the world, especially in countries such as the UK and US where space heating accounts for over 40% of the total residential energy consumption (Department of Energy & Climate Change, 2011). To address this issue, *intelligent* domestic heating systems (IDHS) have been proposed that employ sensors and other technologies so that user-input is minimized while heating is used efficiently by *planning* ahead (i.e., deciding on a set of future actions - for example, turning the heating on ahead of occupancy or turning it off when no one is home). Besides being an active area of research, such IDHS are already starting to make their way into modern homes as commercial products such as the *smart thermostats* from Nest and Honeywell.

In essence, the goal of an IDHS is to compute and execute an *optimal heating plan* to meet some user-defined objective(s). For example, if the users wish to have a warm house on their arrival, then the heating plan is the set of actions (ON/OFF commands to the boiler) to reach the set-point, while minimizing the energy required to do so. At minimum, such planning requires some knowledge of the thermal properties of the home and its heating system. However, these are typically unknown and thus must be learned from historical data in some way (Rogers, Ghosh, Wilcock, & Jennings, 2013). Such learning challenges have been studied in purpose-built commercial buildings and with the onset of an ever increasing instrumentation and cloud computing, the techniques originally developed for commercial buildings are quickly becoming applicable to homes (Rogers et al., 2013). In particular, insights from existing work on adaptive thermal modelling can be adopted for the control of heating in homes in a principled way.

Now, existing works on thermal modelling can be divided into *off-line* and *on-line* learning. In the former case, the thermal model is learned either once or at infrequent intervals and assumed to be fixed over an arbitrary horizon. For example, Gao et al. apply linear regression to historical data in order to learn a fixed thermal model of a room in an office building (Gao & Keshav, 2013). Likewise, Rogers et al. conducted a field-study of 750 UK homes and used non-linear regression to infer the thermal model of each home (Rogers et al., 2013). However, the assumption of a fixed model is inadequate in environments that are highly dynamic, and instead on-line (or *adaptive*) learning model parameters that are updated continuously has been shown to be more resilient and effective in such cases (Kristensen, Madsen, & Jørgensen, 2004). In particular, variants of the extended Kalman filter (EKF) have been used for adaptive thermal modelling in commercial buildings (Fux, Ashouri, Benz, & Guzzella, 2014; Li & Wen, 2014; O'Neill, Narayanan, & Brahme, 2010; Radecki & Hancey, 2012, 2013). Although, such techniques are equally applicable to modern homes, one must consider the challenges peculiar to homes for their effective use.

In more detail, the challenges that set the home apart from other kinds of buildings originate from the fundamental difference in their purposes. This difference manifests itself in terms of the physical structure, appliances, and the occupants' behaviour. For example, commercial buildings usually have standard HVAC (heating, ventilation and air conditioning) systems whose properties are relatively well understood. On the contrary, domestic heating systems are much more diverse in type and, unlike commercial buildings, it is not uncommon for homes to be equipped with multiple heating systems (e.g., an underfloor heating system combined with wall-mounted radiators and/or fan heaters). In addition, household activities such as opening a window or door as winter

retreats or the use of multiple heating systems affect the thermal dynamics of homes. Furthermore, homes typically lack access to accurate local weather information which is a norm for commercial buildings (usually with dedicated roof-top sensors). Similarly, further challenges exist when optimal planning involves factors such as house occupancy which is typically more accurately predictable in commercial buildings than in homes. As such, the challenge of applying robust adaptive thermal models to homes has yet to receive sufficient attention.

To address these shortcomings, we propose an adaptive modelling approach based on stochastic grey-box modelling and the EKF to infer thermal properties of homes. To demonstrate its use, we apply our approach to family home previously used in a well cited smart home experiment called *PreHeat* (Scott et al., 2011). More specifically, we extend the state-of-the-art as follows:

- We provide a general step-by-step formulation for adaptive modelling in homes that uses stochastic grey-box models and the EKF to infer the underlying thermal models.
- We demonstrate the applicability and robustness of our approach by applying it to a family house equipped with underfloor heating in Cambridge, UK. We show via empirical evaluation that the 95th percentile of the absolute prediction error is 0.95°C and 1.37°C for 2 and 4 hours predictions, respectively; in contrast to the corresponding 2.09°C and 3.11°C errors of the existing (historical-average) thermal model.

The rest of the paper is structured as follows. Section 2 reviews the existing work on thermal modelling. Section 3 provides our problem description. Section 4 outlines our challenge and case study. We formulate our thermal models and derive the EKF for our case study in Section 5. Section 6 presents the empirical results of our study while Section 7 concludes and discusses our future work.

2. Related Work

The existing approaches to modelling the thermal interactions of a building can be categorised into forward system identification or *process-driven* and inverse system identification or *data-driven* as follows.

Process-driven approaches use complex system equations that employ thermodynamics principles and knowledge from material science to describe the internal heat processes in a building (Ellis, Hazas, & Scott, 2013). These equations require detailed inputs such as the thermal capacitance, thermal resistance, or thickness of the material (e.g., walls and building envelopes) as well as their dimensions and other relevant properties. This information is generally entered into building information tools that generate building models. These models are then used to simulate the effect of a specific HVAC system on the building. Process-driven approaches have received much attention and there are bespoke software available to facilitate these approaches. Notable examples of such software include the EnergyPlus simulator that is maintained by the US Department of Energy and the TRNSYS simulator from the University of Wisconsin¹. However, the detailed knowledge requirement of process-driven approach is normally feasible only for large commercial buildings where the detailed building construction plans are available.

¹Available at <http://se1.me.wisc.edu/trnsys/> and <http://apps1.eere.energy.gov/buildings/energyplus>

In contrast to the data intensive requirements in process-driven approaches, data-driven approaches utilise the concepts of equivalent thermal parameters (ETPs). These ETPs are used in the system equations to capture the thermodynamics of buildings (Sonderegger, 1978). The number and type of such ETPs are specific to a building, and the task of the modeller is to choose the relevant ETPs based on the prior knowledge of the building (e.g., number of rooms and type of heating systems). This activity is referred to as *grey-box* modelling (see Section 5 for details) and has been used to derive thermal models for buildings for many decades (Andersen, Madsen, & Hansen, 2000). Given a grey-box model, the value of each ETP is usually unknown, and thus, is inferred by applying statistical methods. The exact technique used is selected based on a number of factors such as the complexity (e.g., linear or non-linear) and dynamics of the model (which in turn reflect the complexity of the thermodynamics of the building), available resources (such as processing power and computational memory requirements) and the overall objective of the heating system.

Existing works on data-driven approaches can be divided into off-line and on-line learning, based on how *often* the ETPs are updated. In the former case, the thermal model is learned either once or at infrequent intervals and assumed to be fixed over an arbitrary horizon. For example, Gao et al. applies linear regression to historical data in order to learn a fixed thermal model of a room in an office building (Gao & Keshav, 2013). The heating control system then uses this thermal model to maximise the personal thermal comfort in the workplace. Most recently, the MyJoulo project used specially designed USB loggers to collect data from 750 homes (Rogers et al., 2013). They apply non-linear regression to infer the thermal model of each home in order to provide personalised feedback on energy savings. Similarly, MatchStick uses a non-linear least square method to learn the thermal model of individual rooms in order to predict indoor temperature (Ellis et al., 2013). These predictions are then used to calculate the potential energy savings. In addition, simpler thermal models such as those used in PreHeat employ historical averages to estimate the required parameters (Scott et al., 2011). However, these studies assume a *fixed* model over an arbitrary horizon which is ineffective in highly dynamic environments such as homes, both in terms of the external affect (e.g., weather or seasonal affects) and the occupants' behaviour.

In contrast, adaptive thermal modelling has shown to be a better alternative for buildings where the model is updated frequently using techniques such as Kalman filters and genetic algorithms (Li & Wen, 2014). Although there is no general technique that is useful in all environments, Kalman filters have been used extensively for on-line learning in other domains (Simon, 2006). In particular, Fux et al. applied an EKF to a commercial building (Swiss Alpine Club, Switzerland) for adaptive thermal modelling (Fux et al., 2014). Their EKF models the indoor temperature, leakage rates and solar radiation. They show that their thermal modelling technique is capable of adapting to the seasonal change in the heat flow due to the occupants. Their empirical evaluation shows that their model predicts the indoor air temperature with 0.5°C prediction error for 3 hours and 3°C for 48 hours (2 days). In a similar fashion, other studies have used EKFs for adaptive modelling of buildings to demonstrate prediction of the thermal properties of buildings to their required accuracy (O'Neill et al., 2010; Radecki & Hency, 2012, 2013). Although such studies show the effective use of EKFs for buildings with standard HVAC systems, reliable occupancy schedules and minimal weather effect, the application of EKFs to homes where such conditions do not hold is yet to be explored.

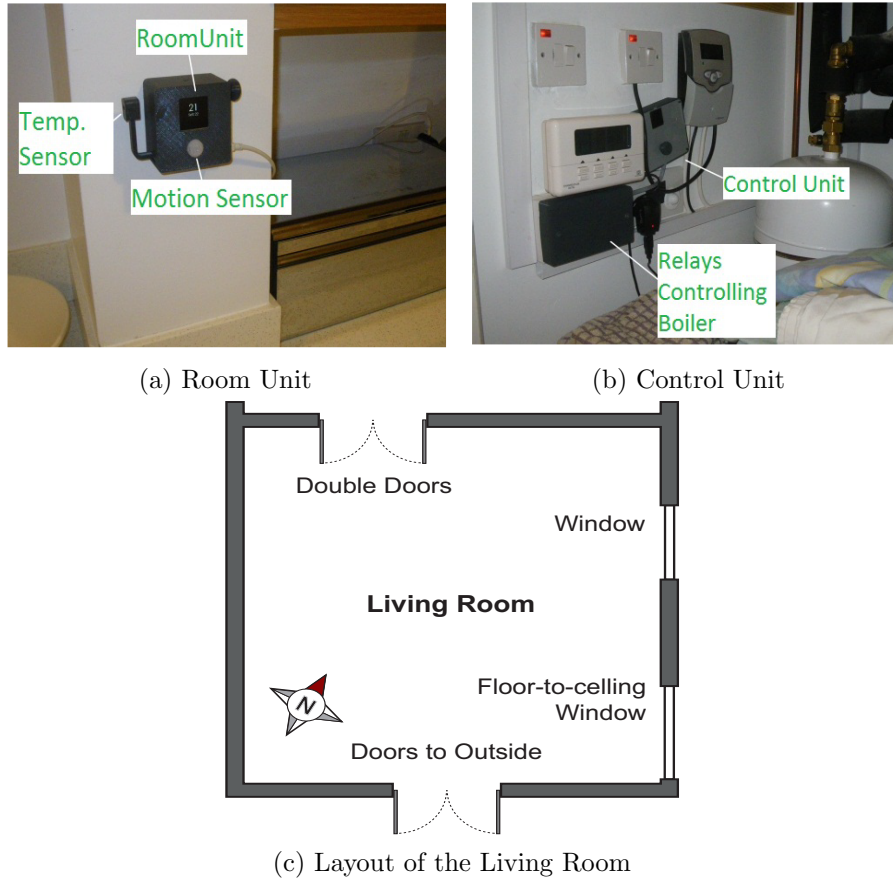


Figure 1.: Case Study - Custom .NET Gadgeteer hardware and the layout of the room.

3. Problem Description

Following our discussion in Sections 1 and 2, our aim is to present an adaptive thermal modelling approach that addresses the common challenges in homes. We note that the exact nature of such challenges may vary between homes depending on a number of factors (e.g., such as the number or type of the heating system(s) in use). However, general challenges arise from a marked absence of detailed knowledge in homes (compared to purpose-built buildings, see Section 1). In particular, we consider the problem of adaptive thermal modelling in homes where (i) the prior knowledge of the physical properties and (ii) the domestic heating system in the home is limited, and (iii) the weather data is not recorded. In order to demonstrate the applicability of our approach, we consider the case of adaptive thermal modelling for a room in a family house in Cambridge as follows.

4. Case Study: Family House In Cambridge

We consider the temperature control of the living room of a family house in Cambridge, UK. The house has both radiators (4 of 10 independently heatable rooms) and underfloor heating (8 of 10 rooms) and is equipped with custom hardware using .NET Gadgeteer to control the heating system and record data. Heating is controlled on a

per-room basis by installing a *room unit* to control wireless radiators valves and a *control unit* to control the underfloor heating valves in a room, as shown in Figure 1a and Figure 1b. The room unit has (i) indoor air temperature (Sensiron SHT15 sensor - accuracy +/- 0.1°C) and (ii) occupancy detection (Panasonic PIR-AMN34111J - passive infra-red motion sensor), while the control unit has no sensors but a relay for controlling the underfloor heating valves. Both units run programmable customer C# software and have 802.15.4 radio modules to establish a wireless mesh communication with a central server PC, where our data is recorded (see (Scott et al., 2011) for more details on this set-up).

We choose the living room for our study as (i) it is largely in use when the house is occupied and (ii) its thermal dynamics are particularly challenging owing to its physical properties and household activities. It has two doors and three windows as shown in Figure 1c and is equipped with underfloor heating (no radiators) and a fan heater that is occasionally used. The use of underfloor heating involves multiple heat transfer processes whereby heat is transferred from the source (i.e., hot-water pipes) to an intermediate thermal mass (i.e., floor) which then slowly leaks to its surroundings (e.g., air, ground, and house envelope). This process introduces thermal lags and additional leakages that must be taken into account, thus making this room interesting from a heating control perspective. In addition, the per-room based heating in the adjacent rooms and the weather outside can affect the indoor and outdoor thermal leakage rates. Furthermore, occupants' activities can have substantial effects on the thermal dynamics of a building and examples of such events in our case include opening of a window-slit in the room as winter retreats or use of the fan heater. Taken together, these factors make this room a good example to consider the challenge of heating control in a family home.

5. Thermal Modelling

In this section, we describe in detail the process of developing thermal models for a home, using our case study as an example. We then show the necessary transformation of a thermal model in order to apply the EKF.

5.1. Formulation of a Thermal Model

We adapt the grey-box modelling approach which uses a combination of the prior physical knowledge of the system and observed data to derive a thermal model. The prior physical knowledge is conveyed by a set of stochastic equations which describe a lumped model of the heat dynamics of the space. This model is then validated against the observed data to measure its consistency and to learn parametrised values that describe the underlying system state (e.g., heater or solar output). An important aspect of the grey-box modelling is the *degree* of its greyness, in other words, its proximity to a white-box model (where sufficient prior knowledge is available to describe system dynamics perfectly) or to a black-box model (where no prior information is available). Intuitively speaking, a model that increasingly relies on more prior knowledge tends to be more accurate but with increasing complexity and at the loss of generality. The task of the grey-box modelling is then to develop a model that strikes a balance between accuracy and complexity within the requirements of the system under consideration. This is generally accomplished by starting out with the simplest feasible model which is then iteratively refined into more complex models to identify the most suitable

	Time		Units
	k	k+1	
Floor temperature	T_f	\hat{T}_f	$^{\circ}\text{C}$
Room air temperature	T_a	\hat{T}_a	$^{\circ}\text{C}$
Envelope temperature	T_e	\hat{T}_e	$^{\circ}\text{C}$
External temperature	T_o	-	$^{\circ}\text{C}$
Heater output	λ_h	$\hat{\lambda}_h$	$^{\circ}\text{C}/\text{hr}$
Global Solar Irradiance	λ_s	-	$\text{J}/\text{m}^2/\text{hr}$
Solar Irradiance (Air)	λ_a	-	$\text{J}/\text{m}^2/\text{hr}$
Solar Irradiance (Envelope)	λ_e	-	$\text{J}/\text{m}^2/\text{hr}$
Heating time ratio	r_h	-	$\in [0, 1]$
Irradiance ratio (Air)	r_a	\hat{r}_a	$^{\circ}\text{C}.\text{m}^2/\text{J}$
Irradiance ratio (Envelope)	r_e	\hat{r}_e	$^{\circ}\text{C}.\text{m}^2/\text{J}$
Leakage rate (Floor, Air)	ϕ_{fa}	$\hat{\phi}_{fa}$	1/hr
Leakage rate (Air, Outside)	ϕ_{ao}	$\hat{\phi}_{ao}$	1/hr
Leakage rate (Air, Envelope)	ϕ_{ae}	$\hat{\phi}_{ae}$	1/hr
Leakage rate (Envelope, Outside)	ϕ_{eo}	$\hat{\phi}_{eo}$	1/hr

Table 1.: Notation used for thermal modelling.

model, as we demonstrate for our case study as follows.

In this case, our prior knowledge of the room includes (i) the use of underfloor heating (ii) binary status of heating (i.e., ON/OFF) at a given time (iii) indoor air temperature and (iv) presence of considerable process noise (due to human activities - see Section 4). It should be noted that there is no knowledge of (i) the heater output (ii) floor temperature (and house *envelope* temperature) (iii) thermal capacities and resistance (i.e., *leakage rates*) of the surroundings and (iv) the room layout (e.g., dimension of the walls or dimension and orientation of the windows). **We note that the notion of leakage rates can be interpreted as the *cumulative* representation of the thermal capacitance and resistance. For example, when the heat transfer between the floor and indoor air is considered then $\phi_{fa} = \frac{1}{R_{fa}C_f}$ where R_{fa} is the thermal resistance between the floor and the indoor air, and C_f is the thermal capacitance of the floor. Similarly, the irradiance ratio for the house envelope can be interpreted as $r_e = \frac{A_e}{C_e}$ where A_e is the effective area of the house envelope that receives irradiance and C_e is its thermal capacitance. This alternative interpretation is useful to build the RC-network of a thermal model.** Now, the variables (or system states) are considered *hidden* or *non-observable* and therefore must be inferred indirectly, where required. Table 1 introduces our notation for system states. Our convention is to use a X for scalars, \mathbf{X} for vectors and matrices at time period k and $\hat{\mathbf{X}}$ for estimates at $k + 1$. The equations assume the time step to be 1 hour for clarity, and all the rates are scaled appropriately if the time step is different. In the following paragraphs, we outline a number of increasingly complex thermal models.

Single-Box Model (T_a): The single-box model (T_a) is the simplest feasible model in our study where the transfer of heat is assumed to be direct between the heater and

the indoor air which then leaks to the outside. Here, no thermal delay for the heat transfer (between the heater and indoor air) exists, enabling us to measure the effect of thermal delays in other models. The basic model is described as:

$$\hat{T}_a = T_a + r_h \lambda_h + \phi_{ao}(T_o - T_a) \quad (1)$$

Underfloor Heating Model ($T_a T_f$): Here, we model the transfer of heat from the heater to the indoor air via an intermediary thermal mass (i.e., the floor), thus introducing a thermal lag in the heat transfer process. The indoor air then leaks heat to the outside, as follows:

$$\hat{T}_f = T_f + r_h \lambda_h + \phi_{fa}(T_a - T_f) \quad (2)$$

$$\hat{T}_a = T_a + \phi_{fa}(T_f - T_a) + \phi_{ao}(T_o - T_a) \quad (3)$$

Underfloor and Envelope Model ($T_a T_f T_e$): This is an extension of the underfloor heating model $T_f T_a$ where the heat escapes to outside via the house envelope or structure.

$$\hat{T}_f = T_f + r_h \lambda_h + \phi_{fa}(T_a - T_f) \quad (4)$$

$$\hat{T}_a = T_a + \phi_{fa}(T_f - T_a) + \phi_{ae}(T_e - T_a) \quad (5)$$

$$\hat{T}_e = T_e + \phi_{ae}(T_a - T_e) + \phi_{eo}(T_o - T_e) \quad (6)$$

Complete Model ($T_a T_f T_e S$): The complete model extends the underfloor and envelope model $T_f T_a T_e$ to include the effect solar irradiation on the indoor temperature as well as on the the house envelope.

$$\hat{T}_f = T_f + r_h \lambda_h + \phi_{fa}(T_a - T_f) \quad (7)$$

$$\hat{T}_a = T_a + r_a \lambda_s + \phi_{fa}(T_f - T_a) + \phi_{ae}(T_e - T_a) \quad (8)$$

$$\hat{T}_e = T_e + r_e \lambda_s + \phi_{ae}(T_a - T_e) + \phi_{eo}(T_o - T_e) \quad (9)$$

Having outlined our thermal models, we next discuss and derive the EKF.

5.2. Applying the Extended Kalman Filter for the Complete Model

The extended Kalman filter (EKF) applies the Kalman filter to systems where the state transition and observation models can be non-linear functions of the system state (Simon, 2006). The EKF has been used extensively in numerous applications (including heating systems) and is considered to be the de-facto standard in many domains where non-linear estimation is required (Grewal & Andrews, 2011). The EKF is described by a set of standard equations that operate on a stochastic linear state-space model. Therefore, it is necessary to transform our thermal model to this representation before the EKF can be applied. Here, we show the transformation of our complete model ($T_a T_f T_e S$) only. We then apply the EKF for non-linear parameter estimation in this model. We note that the EKF for underfloor heating systems in general, and for our other models in particular, can be derived by excluding the irrelevant parameters from $T_a T_f T_e S$. Similarly, $T_a T_f T_e S$ can be expanded to introduce more parameters (e.g., heat transfers from adjacent rooms) as the heat exchanges in a home are better understood. Now, we first turn $T_a T_f T_e S$ to a set of stochastic differential equations (SDEs) (for

details on applying SDEs to thermal systems, see (Andersen et al., 2000; Bacher & Madsen, 2011)), as follows:

$$\hat{T}_f = T_f + r_h \lambda_h + \phi_{fa}(T_a - T_f) + E_f \quad (10)$$

$$\hat{T}_a = T_a + r_a \lambda_s + \phi_{fa}(T_f - T_a) + \phi_{ae}(T_e - T_a) + E_a \quad (11)$$

$$\hat{T}_e = T_e + r_e \lambda_s + \phi_{ae}(T_a - T_e) + \phi_{eo}(T_o - T_e) + E_e \quad (12)$$

E_f , E_a and E_e represent a *white* process noise (see Matrix \mathbf{Q} in this section). Now, we need four matrices to convert above Equations (10,11 and 12) into state space representation as follows.

State Space Vector (\mathbf{X}) consists of state variables that sufficiently describe the state of the system at any given time. For our complete model $T_a T_f T_e S$, its state variables can be represented as a vector:

$$\mathbf{X} = [\lambda_h \quad r_a \quad r_e \quad \phi_{fa} \quad \phi_{ae} \quad \phi_{eo} \quad T_f \quad T_a \quad T_e]^T$$

State Transition Matrix (\mathbf{F}) describes the evolution of all state variables with the system transition. For the complete model $T_a T_f T_e S$, we formulate its state transition matrix as:

$$\mathbf{F} = \begin{bmatrix} 1 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 1 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 1 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 1 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 1 & 0 & 0 & 0 \\ r_h & 0 & 0 & T_a - T_f & 0 & 0 & 1 & 0 & 0 \\ 0 & \lambda_s & 0 & T_f - T_a & T_e - T_a & 0 & 0 & 1 & 0 \\ 0 & 0 & \lambda_s & 0 & T_a - T_e & T_o - T_e & 0 & 0 & 1 \end{bmatrix}$$

Jacobian Transition Matrix ($d\mathbf{F}$) is the matrix of all first-order partial derivatives of \mathbf{X} . To compute $d\mathbf{F}$, we decompose \mathbf{X} into its *components functions* as follows (see (Simon, 2006) for details). Let $f_i(\mathbf{X}) = \hat{x}_i \mid \hat{x}_i \in \hat{\mathbf{X}}$ define the component functions as follows.

$$\begin{bmatrix} f_1 \\ f_2 \\ f_3 \\ f_4 \\ f_5 \\ f_6 \\ f_7 \\ f_8 \\ f_9 \end{bmatrix} = \begin{bmatrix} \lambda_h \\ r_a \\ r_s \\ \phi_{fa} \\ \phi_{ae} \\ \phi_{eo} \\ T_f + r_h \lambda_h + \phi_{fa}(T_a - T_f) \\ T_a + r_a \lambda_s + \phi_{fa}(T_f - T_a) + \phi_{ae}(T_e - T_a) \\ T_e + r_e \lambda_s + \phi_{ae}(T_a - T_e) + \phi_{eo}(T_o - T_e) \end{bmatrix} \quad (13)$$

Now, the Jacobian matrix $d\mathbf{F}$ can be computed with respect to the partial derivatives of the component functions as:

$$dF = \begin{bmatrix} \frac{\partial f_1}{\partial \lambda_h} & \frac{\partial f_1}{\partial r_a} & \frac{\partial f_1}{\partial r_e} & \frac{\partial f_1}{\partial \phi_{af}} & \frac{\partial f_1}{\partial \phi_{ae}} & \frac{\partial f_1}{\partial \phi_{eo}} & \frac{\partial f_1}{\partial T_f} & \frac{\partial f_1}{\partial T_a} & \frac{\partial f_1}{\partial T_o} \\ \frac{\partial f_2}{\partial \lambda_h} & \dots & \dots & \dots & \dots & \dots & \dots & \dots & \frac{\partial f_2}{\partial T_o} \\ \vdots & \vdots & \vdots & \vdots & \vdots & \vdots & \vdots & \vdots & \vdots \\ \frac{\partial f_9}{\partial \lambda_h} & \dots & \dots & \dots & \dots & \dots & \dots & \dots & \frac{\partial f_9}{\partial T_o} \end{bmatrix}$$

$$dF = \begin{bmatrix} 1 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 1 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 1 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 1 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 1 & 0 & 0 & 0 \\ r_h & 0 & 0 & \frac{\partial f_7}{\partial \phi_{fa}} & 0 & 0 & \frac{\partial f_7}{\partial T_f} & \phi_{fa} & 0 \\ 0 & \lambda_s & 0 & \frac{\partial f_8}{\partial \phi_{fa}} & \frac{\partial f_8}{\partial \phi_{ae}} & 0 & \phi_{fa} & \frac{\partial f_8}{\partial T_f} & \phi_{ae} \\ 0 & 0 & \lambda_s & 0 & \frac{\partial f_9}{\partial \phi_{ae}} & \frac{\partial f_9}{\partial \phi_{eo}} & 0 & \phi_{ae} & \frac{\partial f_9}{\partial T_o} \end{bmatrix}$$

where

$$\begin{aligned} \frac{\partial f_7}{\partial \phi_{fa}} &= T_a - T_f \quad , \quad \frac{\partial f_7}{\partial T_f} = 1 - \phi_{fa} \\ \frac{\partial f_8}{\partial \phi_{fa}} &= T_f - T_a \quad , \quad \frac{\partial f_8}{\partial \phi_{ae}} = T_e - T_a \\ \frac{\partial f_8}{\partial T_a} &= 1 - \phi_{fa} - \phi_{ae} \quad , \quad \frac{\partial f_9}{\partial \phi_{ae}} = T_a - T_e \\ \frac{\partial f_9}{\partial \phi_{eo}} &= T_o - T_e \quad , \quad \frac{\partial f_9}{\partial \phi_{ee}} = T_a - T_o \end{aligned}$$

Process Noise Matrix (\mathbf{Q}) defines the process noise for each parameter. We define it as (i) *static* for parameters other than temperatures and as (ii) *dependent* (on the related random variables) noise for temperatures. For example, the process noise in T_f is dependent on the heater output λ_h and leakage ϕ_{fa} so that as either of these variable increase, so does the process noise in T_f . Intuitively, the dependent noise here indicates that we expect T_f to be more *noisy*, when the heater is on than when it is off. Thus, the process noise matrices is given by:

$$\mathbf{Q} = \begin{bmatrix} Q_h & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & Q_s & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & Q_s & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & Q_\phi & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & Q_\phi & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & Q_\phi & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & E_f & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & E_a & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & E_e \end{bmatrix}$$

where each diagonal entry is the process noise to corresponding state in \mathbf{X}^T , R is the

measurement noise, and

$$\begin{aligned} E_f &= R(\lambda_h^2 + \phi_{fa}^2) \\ E_a &= R(r_a + \phi_{fa}^2 + \phi_{ae}^2) \\ E_e &= R(r_e + \phi_{ae}^2 + \phi_{eo}^2) \end{aligned}$$

This transforms our $T_a T_f T_e S$ into an equivalent stochastic state-space representation where the EKF can be applied. Now, the standard EKF equations can be used to calculate the posteriori distribution (i.e., *posteriori* estimate and covariance) and residuals as follows.

Priori Estimate	$\hat{\mathbf{X}}^- = \mathbf{F}\mathbf{X}$
Priori Error Covariance	$\hat{\mathbf{P}}^- = \mathbf{d}\mathbf{F}\mathbf{P}\mathbf{d}\mathbf{F}^T + \mathbf{Q}$
Residual	$Y = T_a - \mathbf{H}\mathbf{X}^-$
Residual Covariance	$S = \mathbf{H}\hat{\mathbf{P}}^-\mathbf{H}^T + R$
Kalman Gain	$\mathbf{K} = (\hat{\mathbf{P}}^-\mathbf{H}^T)/S$
Posteriori Estimate	$\hat{\mathbf{X}}^+ = \hat{\mathbf{X}}^- + \mathbf{K}Y$
Posteriori Error Covariance	$\hat{\mathbf{P}}^+ = (\mathbf{I} - \mathbf{K}\mathbf{H})\hat{\mathbf{P}}^-$

Here, $\hat{\mathbf{X}}^+$ contains the posteriori estimates for $T_a T_f T_e S$ which we show in Section 6.2 for a given dataset.

6. Empirical Evaluation

Having outlined our adaptive thermal modelling and the EKF formation, we now present an empirical evaluation of our work that demonstrate its advantages and potential use for heating control in homes. Our objectives are to (i) show how our model selection is driven towards a suitable thermal model and (ii) present a systematic analysis of our estimated parameters to show their accordance with reality.

6.1. Data Collection

We use temperature data from November, 2011 to March, 2012 (150 days in total) that was collected at 0.1Hz using the custom hardware as detailed in Section 4. This duration of our evaluation period is selected to investigate the performance of our adaptive thermal model for infrequent events (e.g., turning off the underfloor heating for holidays and the occasional use of the fan heater) as well as the effect of seasonal changes.

Using this data, for a given time period (i.e., estimation period - here it is 5 minutes for the EKF), we calculate its heating ratio r_h (i.e, how many minutes the heating was on in that particular time period) and the air temperature T_a in the living room

³We use linear regression to estimate T_a or T_o where needed.

⁴<https://www.cl.cam.ac.uk/research/dtg/weather/>

⁵<http://re.jrc.ec.europa.eu/pvgis/apps4/pvest.php>. We note that finer resolution solar irradiance can improve prediction accuracy in Section 6.2.

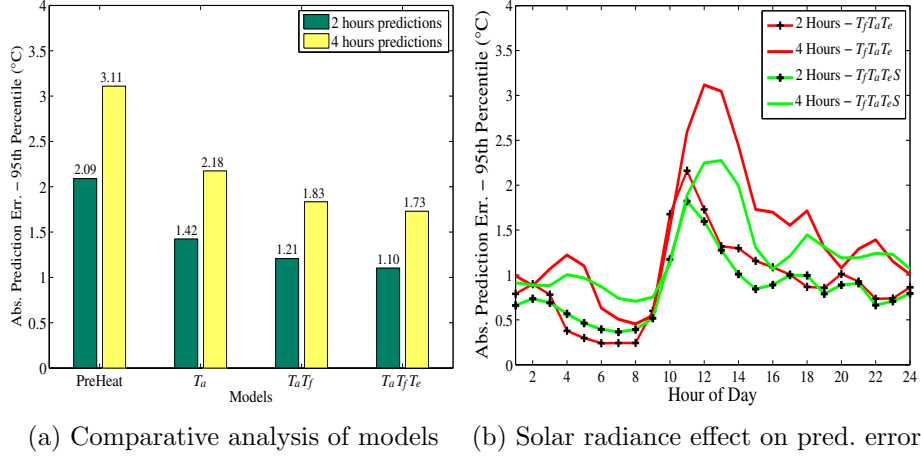


Figure 2.: Prediction error and model comparison analysis.

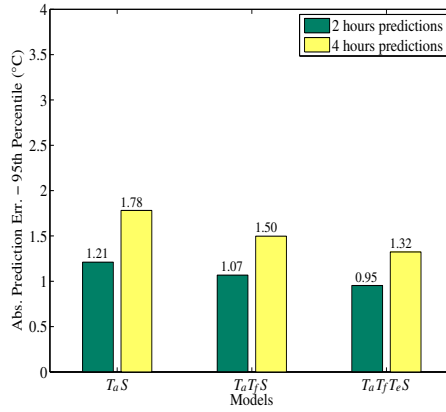


Figure 3.: Thermal models with solar irradiance

at the end of this time period. As we do not consider the possibility of the on-site collection of weather data at home, we use publicly available weather data³. More specifically, to calculate the external temperature T_o , we use data from the University of Cambridge Computer Laboratory (30 minutes interval) due to its proximity and reliability⁴. In addition, to estimate the solar irradiance λ_s received in an estimation period, we use data from the EU Joint Research Commission due to its reliability and general coverage of Europe and Africa⁵. This dataset consists of solar irradiance on a typical day for a given month (note that this means that all days in any given month have the same solar irradiance). This data consists of average solar irradiance for 15 minutes interval and we assume this to be constant within this interval.

6.2. Evaluating the Thermal Models

The underfloor heating systems are associated with significant thermal delays and typically affect the heat process at a much slower rate, compared to other heating systems (e.g., wall-mounted radiators and forced-air systems). Generally, they can take up to several hours to reach the set point, and even longer if the heating is turned off for longer periods of time (e.g., when the household is away for holidays). Systems with such thermal lags require the heating control to be able to plan and predict

hours ahead. While the time and accuracy requirements of prediction vary based on the system and its objectives, we find that in our case, predicting 2 and 4 hours ahead with the respective accuracy of $\pm 1^\circ\text{C}$ and $\pm 1.5^\circ\text{C}$, is sufficient for an effective heating control. These accuracy requirements are in line with the existing user comfort models in literature (Section 2). We note that at present, the PreHeat control system assumes a constant $\lambda_h = 1.5^\circ\text{C}/\text{hour}$ heating output (derived from historical data) in the living room when the underfloor heating is on (Scott et al., 2011). We use this as our baseline to find out if our adaptive models can offer any improvements over this static method to justify their use.

In more detail, to show the comparative performance of the EKF models over the existing PreHeat, **we evaluate all models using our collected data in Section 6.2 and analyse their absolute errors (calculated as the absolute difference of the actual and predicted temperature) for 2 and 4 hours predictions within 95th percentile, as shown in Figure 2a.** For comparison, the 95th percentiles error for PreHeat predictions is 2.09°C and 3.11°C , respectively (note that the 95th percentile roughly corresponds to the confidence interval of 2 standard deviations). We note that even the basic T_a model offers improvement over the existing PreHeat, confirming that stochastic thermal modelling is useful in this particular problem. It is also evident that as modelling increasingly captures the physical properties of the building, the prediction accuracy increases further. More specifically, as we incorporate the thermal mass of the floor (T_aT_f) and the house envelope ($T_aT_fT_e$), our prediction accuracy increases. However, it should be noted that this accuracy comes at the cost of modelling additional state variables, and thus the need for more training data.


To drive our model selection further, we analyse the nature of the prediction error with $T_aT_fT_e$. In particular, Figure 2b shows the absolute prediction errors (95th percentile) of $T_aT_fT_e$ against the hour of day. We note that the prediction error is relatively large during the daytime hours, hinting at the importance of solar irradiance in the heat dynamics of the room. This is indeed the case as shown in Figure 2b where the prediction accuracy improves when the solar irradiance is factored in $T_aT_fT_eS$. However, this correlation, though reduced, is not completely eliminated in $T_aT_fT_eS$. This is due to the fact that our solar irradiance data consists of monthly averages only and therefore, can only be used as an approximation to the irradiance on actual days.

Although Figure 2b confirms that solar irradiance can make predictions more accurate, it does not indicate which of our three models will benefit the most by factoring in the solar irradiance term. To explore this effect, we include solar irradiance in all models and calculate the 95th percentiles of prediction error for every model, as shown in Figure 3. It is evident that $T_aT_fT_eS$ is the most accurate model. We also note that the prediction error is 0.95°C which meets our accuracy requirements. We note that this accuracy is expected to improve with the inclusion of more frequent and finer-grained external temperature and solar irradiance data.

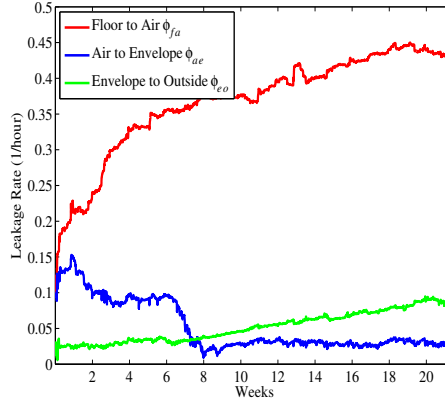
6.3. *Interpreting the Learned Estimates*

As we showed in Section 6.2, the complete model $T_aT_fT_eS$ best describes our data and the system dynamics in our case study. However, a recommended practice is to analyse the learned parameters as it provides (i) means to check whether the learned parameters match with reality and other prior knowledge that cannot be factored in the model in a principled way (e.g., a subjective comparison of two unknown leakage

⁶see <http://www.telegraph.co.uk/property/3345024/Will-you-really-feel-the-heat-under-your-feet.html>



(a) Estimated Temperatures (b) Estimated Heater and Solar output



(c) Estimated leakage rates

Figure 4.: Interpreting the Learned Parameters

rates), and (ii) useful insights into the dynamics of the building that could be used for further model selection, if needed (Bacher & Madsen, 2011). In line with this recommendation, our analysis of the estimates of the $T_a T_f T_e S$ is as follows.

Figure 4a shows our estimates for the floor temperature T_f , the envelope temperature T_e , and the observed external temperature T_o for our simulation period. We notice that T_f is below the cut-off temperature of 29°C (note, this is the surface temperature) for solid floors in a typical underfloor heating system⁶. We also notice that when the heating is turned off for longer periods of time (Weeks 5, 11 and 13), the estimated T_f drops to almost 16°C. In such cases, T_f takes several hours (7 hours for Week-5 and 5 hours for others) to reach a comfortable temperature, and whereupon the long-term (4 hours) predictions become essential for optimal control. We also notice that T_e closely follows T_o which indicates that the temperature of the house envelope has a strong correlation with the external temperature.

Figure 4b shows the estimated heater output λ_h and solar irradiance λ_a and λ_e that is received in the living room and the house envelope, respectively. Here, $\lambda_a = r_a \lambda_s$ and $\lambda_e = r_e \lambda_s$ (Section 5). We observe that λ_h , λ_a and λ_e have notable fluctuations in the first two weeks, which start decreasing in the following weeks. This indicates that these estimates are learned at a slower rate (compared to say T_f or T_e). However, it should be noted that there are fewer times when λ_h , λ_a and λ_e are at work (e.g., heating is on only for a few hours a day) and therefore, it takes longer to accurately

estimate these parameters. We also note that some initial values do not match reality. In particular, λ_a sometimes indicates a *negative* solar irradiance during the early stages of learning which certainly cannot be true. The heater output λ_h also varies which, ignoring the fluctuations in the first week, can be understood to account for the *missing* factors such as the heating received from the surroundings (including the adjacent and above rooms which both have the underfloor heating and radiators), seldom use of the fan heater in the living room and other human activities. Overall, the heater output averages to slightly over $\pm 1.5^\circ\text{C}/\text{hour}$ which confirms our observation of the historical mean used in the PreHeat system.

Figure 4c shows the leakage rates ϕ_{fa} (floor to air), ϕ_{ae} (air to envelop) and ϕ_{eo} (envelop to outside). We notice that ϕ_{fa} is significantly greater than ϕ_{ae} and ϕ_{eo} , which confirms that the heat transfer is easier between the floor and the air. We also notice that ϕ_{fa} has a general upward trend that correlates with T_o . While it is difficult to understand the relation between the two, it becomes clear when the ground temperature is considered. Although our model does not consider leakage to the ground, in reality, the heat from the floor leaks towards both the ground and the indoor air. The seasonal change in the ground and external air temperature has a positive correlation and as the ground temperature increases, the amount of heat escaping from the floor to the ground decreases and thus more heat escapes to the indoor air. This indirect relation makes ϕ_{fa} and T_o negatively correlated and, in the absence of the ground temperature, $T_a T_f T_e S$ explains this phenomena as if the floor leakage rate changes. More specifically, ϕ_{fa} has a general upward trend in Week-2 to 7 and Week-11 to 14 when T_o have the opposite trend, and vice versa for Week-18 to 22. It should be noted that the ground temperature is generally much more stable than the outside air temperature and sudden drops in T_o do not affect the ground temperature significantly (e.g., Day-3 in Week-14). Overall, the observation of the relation between ϕ_{fa} and T_o gives us an insight that extending $T_a T_f T_e S$ to include the ground temperature, could lead to better estimates for ϕ_{fa} and thus improved predictions. In addition, the lower values of ϕ_{ae} and ϕ_{eo} can also be interpreted as an indication of good insulation in the house. We also notice that ϕ_{ae} has a downward trend over Week-1 to 8, while ϕ_{eo} appears to be relatively stable. However, it should be noted that our prior beliefs for ϕ_{ae} and ϕ_{eo} are different and the trend appears to be in the correction period (i.e., to estimate ϕ_{ae} correctly) until Week-8 and the variation in ϕ_{ae} afterwards is much smaller.

7. Conclusions and Future Work

Learning a thermal model of the home is an essential task for IDHS for optimal control of heating. Adaptive thermal modelling techniques such as EKF's that previously has been used in purpose-built buildings can also be used in homes with necessary modification to tackle the challenges unique to homes. In this study, we provide a general step-by-step formulation of adaptive thermal modelling in homes based on the stochastic thermal models. We note that this formulation is *general* in the sense that it describes a generic method whereby (i) the thermo-dynamics in a home can be converted to a stochastic state-space model and (ii) an EKF is derived for estimations of the ETPs. The thermal models in general, and the ones we discussed, are specific

⁷An example is our previously developed Smart Home Framework that enables rapid modelling of smart homes and communities (Alam, Alan, Rogers, & Ramchurn, 2013). Available online at <http://www.smarthomeframework.ecs.soton.ac.uk/>

to the home under consideration and therefore should be altered as needed.

In addition, the challenges of applying adaptive modelling in homes vary depending on a number of factors (see Section 3), however, homes typically lack the information that are available in large commercial buildings. Following this, our case study assumes very limited prior knowledge (see Section 4) and uses publicly available weather data. We demonstrate via empirical evaluation that despite such challenges in homes, our EKF-based adaptive modelling approach is useful in estimating the unknown ETPs but also robust in the presence of considerable process noise and seasonable changes (Section 6). More specifically, we show that the 95th percentile of the absolute prediction error (of our complete model $T_a T_f T_e S$) is 0.95°C and 1.37°C for 2 and 4 hours predictions, respectively.

It is also evident that our adaptive thermal modelling approach has the potential to capture the thermal dynamics of a house in a sufficient (i.e., by incorporating increasingly complex thermal models) and accurate (e.g., accurate temperature prediction) manner. These properties, coupled with the adaptive nature of our approach, are crucial to the efficient control of heating system and therefore, have wider implications in improving energy performance in homes and thermal comfort. For example, the shorter feedback duration of an adaptive systems (compared to the fixed models, see Section 2) means that any additional leakages (e.g., opened windows) or a newly introduced secondary heating source can be quickly adjusted. This can lead to more efficient control of space-heating as well as an active management of thermal comfort. Similarly, short terms weather changes (e.g., cold spills or rains) can affect the heat dynamics of a house (e.g., leakage rates) and, as evident in our empirical evaluations, an adaptive approach is better at responding at such events. In contrast, with fixed thermal models, the heating system response is dependent on the update frequency that can vary from days to weeks, leading to negative impacts on energy use and users thermal comfort.

We also note that the detailed thermal analysis we presented in Figure 3 and 4 can be used to understand occupants behavioural patterns in general. Such insights can be used to improve energy performance of homes or thermal comfort. For example, identification of high leakage rates in particular rooms may indicate poor insulation in a house (or part thereof). Similarly, rapid changes in leakage rates may indicate behaviours such as window opening (e.g., when they feel warm) and this information can be used to infer thermal comfort level of users.

For our future work, our immediate goals are to (i) deploy our adaptive thermal modelling module (written in C#) in selected homes equipped with the .NET Gagde-teer hardware and PreHeat system in selected homes, in order to demonstrate the effective use our approach across a range of domestic heating systems, and (ii) release our code as an open-source framework for adaptive thermal modelling to researchers to enable rapid development of thermal models, much in the same way as open-source rapid development frameworks are available for other energy domains⁷.

8. Acknowledgements

University of Southampton was supported by the EPSRC grant EP/K503770/1 under the Impact Acceleration Account. The authors are thankful to Steven Reece, Univer-

sity of Oxford, for discussions on the use of the EKF.

References

- Alam, M., Alan, A., Rogers, A., & Ramchurn, S. D. (2013, September). Towards a smart home framework. *5th ACM Workshop On Embedded Systems For Energy-Efficient Buildings (BuildSys2013)*. Retrieved from <http://eprints.soton.ac.uk/357187/>
- Andersen, K. K., Madsen, H., & Hansen, L. H. (2000). Modelling the heat dynamics of a building using stochastic differential equations. *Energy and Buildings*, *31*(1), 13–24.
- Bacher, P., & Madsen, H. (2011). Identifying suitable models for the heat dynamics of buildings. *Energy and Buildings*, *43*(7), 1511 - 1522. Retrieved from <http://www.sciencedirect.com/science/article/pii/S0378778811000491> doi:
- Department of Energy & Climate Change. (2011). *Estimates of heat use in the United Kingdom in 2011* (Tech. Rep.). Department of Energy & Climate Change.
- Ellis, C., Hazas, M., & Scott, J. (2013, April). Matchstick: A room-to-room thermal model for predicting indoor temperature from wireless sensor data. In *Proceedings of information processing in sensor networks 2013*. ACM. Retrieved from <http://research.microsoft.com/apps/pubs/default.aspx?id=184179>
- Fux, S. F., Ashouri, A., Benz, M. J., & Guzzella, L. (2014). EKF based self-adaptive thermal model for a passive house. *Energy and Buildings*, *68*, Part C(0), 811 - 817. Retrieved from <http://www.sciencedirect.com/science/article/pii/S0378778812003039> doi:
- Gao, P. X., & Keshav, S. (2013). Optimal Personal Comfort Management Using SPOT+. In *Proceedings of the fifth acm workshop on embedded systems for energy-efficient buildings* (pp. 1–8).
- Grewal, M. S., & Andrews, A. P. (2011). *Kalman filtering: Theory and practice using matlab*. John Wiley & Sons.
- Kristensen, N. R., Madsen, H., & Jørgensen, S. B. (2004). Parameter estimation in stochastic grey-box models. *Automatica*, *40*(2), 225–237.
- Li, X., & Wen, J. (2014). Review of building energy modeling for control and operation. *Renewable and Sustainable Energy Reviews*, *37*(0), 517 - 537. Retrieved from <http://www.sciencedirect.com/science/article/pii/S1364032114003815> doi:
- O’Neill, Z., Narayanan, S., & Brahme, R. (2010). Model-based thermal load estimation in buildings. In *Proceedings of simbuild 2010, the fourth national conference of ibpsa-usa* (pp. 474–481).
- Radecki, P., & Hency, B. (2012). Online building thermal parameter estimation via unscented Kalman filtering. In *American control conference (acc), 2012* (pp. 3056–3062).
- Radecki, P., & Hency, B. (2013). Online thermal estimation, control, and self-excitation of buildings. In *Decision and control (cdc), 2013 ieee 52nd annual conference on* (pp. 4802–4807).
- Rogers, A., Ghosh, S., Wilcock, R., & Jennings, N. R. (2013). A scalable low-cost solution to provide personalised home heating advice to households. In *Proceedings of the fifth acm workshop on embedded systems for energy-efficient buildings* (pp. 1–8).
- Scott, J., Bernheim Brush, A., Krumm, J., Meyers, B., Hazas, M., Hodges, S., & Villar, N. (2011). Preheat: Controlling home heating using occupancy prediction. In *Proceedings of the thirteenth acm international conference on ubiquitous computing* (pp. 281–290).
- Simon, D. (2006). *Optimal state estimation: Kalman, h infinity, and nonlinear approaches*. John Wiley & Sons.
- Sonderregger, R. (1978). Diagnostic test determining the thermal response of a house. *ASHRAE Transactions*, *84*, Part 1, 691–702.