

# AdaHeat: A General Adaptive Intelligent Agent for Domestic Heating Control

Athanasios Aris Panagopoulos, Muddasser Alam, Alex Rogers, Nicholas R. Jennings  
School of Electronics and Computer Science  
University of Southampton, UK  
{ap1e13,moody,acr,nrj}@ecs.soton.ac.uk

## ABSTRACT

Improving the energy efficiency of domestic heating systems can lead to a major reduction in energy consumption and the corresponding CO<sub>2</sub> emissions. To this end, intelligent domestic heating agents (IDHAs) aim to operate domestic heating systems more efficiently with minimum user input. In this work, we propose a new general IDHA that balances heating cost and thermal discomfort in an infinite horizon optimization manner, learns an adaptive thermal model of the system under control on-line and plans a heating schedule that fully exploits the probabilistic occupancy estimates. Importantly, our agent adapts to the user preferences in balancing heating cost and thermal discomfort, as it relies on a single parametrization variable that is learned on-line, and is able to consider a wide range of heating systems typically employed in domestic settings. The backbone of our IDHA is an adaptive model predictive control approach along with a new general planning algorithm that utilizes dynamic programming. We present a thorough evaluation of our approach, and show its effectiveness in terms of Pareto efficiency and usability criteria against state-of-the-art IDHAs. By so doing, we also conduct a comprehensive characterization of existing IDHAs to provide significant insights about their performance in different operational settings.

## Categories and Subject Descriptors

I.2.8 [Artificial Intelligence]: Problem Solving, Control Methods, and Search—*Control theory, Dynamic programming, Graph and tree search strategies*

## General Terms

Algorithms, Performance, Reliability, Experimentation

## Keywords

Machine Learning; Control; Energy Savings

## 1. INTRODUCTION

In many countries, such as the UK and the US, the domestic sector accounts for more than 20% of the total energy consumption, and over 40% of this share is related to space

heating.<sup>1</sup> As such, improving the energy efficiency of domestic heating systems can lead to a major reduction in energy consumption and the corresponding CO<sub>2</sub> emissions. To this end, intelligent domestic heating agents (IDHAs) aim to operate such systems (i.e., optimize the heating control process) more efficiently than current manual (programmable or static) thermostat control, with minimum user input [10].

Now, the goal of any heating automation system is to balance heating cost and the occupant’s thermal discomfort according to their preferences—this balancing considers a non-trivial bi-objective optimization task. In this context, energy research has long been preoccupied with developing such supervisory control systems for *non-domestic* buildings (e.g., [10, 13]). However, more recently, with the onset of ever-increasing house instrumentation and cloud computing, experimental IDHAs are also starting to emerge (e.g., [24, 28, 21, 31, 29]) and have already made their way into modern homes as commercial products (e.g., Nest, Honeywell and Hive). Such autonomous agents become essential in domestic heating settings as the latter provide additional challenges over their non-domestic counterparts.

In particular, the thermal dynamics of domestic buildings are harder to model accurately than their non-domestic counterparts as: (i) the occupant’s activity is more diverse and highly affects the thermal dynamics of the house (e.g., opening a window, operating an auxiliary heater, or cooking) [20, 12]; (ii) the temperature in adjacent buildings or rooms is rarely observed and/or predicted [20, 17]; and (iii) the local weather observations and forecasting reports are usually less accurate due to lack of appropriate instrumentation [20, 9]. In addition, the occupancy schedule—which is an essential input to any thermal comfort model (as any comfort is experienced only when the space is occupied)—is typically unknown in domestic settings and needs to be predicted [19]—in contrast to commercial buildings where it is typically fixed. In this context, all proposed predictive approaches inevitably retain an uncertainty over this schedule which is modeled in the form of probabilistic estimates [19]. In the presence of this uncertainty, sacrificing thermal comfort is typically inevitable to avoid extreme heating cost. As such, dealing with this uncertainty and matching the occupant’s preferences in balancing discomfort and cost arises as a significant challenge in IDHAs. This is exacerbated by the fact that even for a single household these preferences vary over time as they are affected by a range of time-varying factors (e.g., availability of money; health conditions). Finally, domestic heating systems are much more diverse than

**Appears in:** *Proceedings of the 14th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2015), Bordini, Elkind, Weiss, Yolum (eds.), May 4–8, 2015, Istanbul, Turkey.*  
Copyright © 2015, International Foundation for Autonomous Agents and Multiagent Systems (www.ifaamas.org). All rights reserved.

<sup>1</sup>Based on data from [www.gov.uk](http://www.gov.uk) and [www.eia.gov](http://www.eia.gov).

those used in non-domestic buildings which calls for a general IDHA (that is able to handle a variety of them).

That said, a number of IDHAs have been proposed in the literature (e.g., [24, 28, 21, 31, 29]). However, they typically suffer from several drawbacks: (i) they usually rely on a simple experimental thermal model which is not reliable in practice and suitable only for proof-of-concept systems; if not (ii) they do not deal with the highly dynamic nature of house thermal characteristics; (iii) they do not provide a way of choosing the parameterizable coefficients in balancing heating cost and thermal discomfort—the important challenge of matching the occupant’s preferences is usually disregarded in IDHAs; (iv) they usually rely on heuristic control approaches in dealing with occupancy uncertainty (without providing any guarantees or intuition regarding the performance loss from an approach that fully exploits the probabilistic estimates); if not (v) they rely on computationally expensive approaches that limit their applicability only to experimental settings; and (vi) they are usually heating-system-specific. In addition to the above limitations, there is also a lack of comparison among IDHAs, as those are usually benchmarked against simple static timer programs such as “always-on” or “pre-scheduled” heating.

To address these shortcomings, we propose a new general IDHA, AdaHeat, that balances heating cost and thermal discomfort in an infinite horizon optimization manner, learns an adaptive thermal model of the system under control on-line and does planning to fully exploit the occupancy probabilities. To this end, our agent employs a model predictive control (MPC) approach utilizing adaptive gray-box thermal modeling and a new general algorithm for planning that fully exploits the probabilistic occupancy estimates via dynamic programming. As such AdaHeat: (i) is able to effectively account for the highly dynamic thermal characteristics of houses, (ii) is able to work in conjunction with both linear and non-linear optimization objectives and system models, (iii) and is general enough to consider a wide range of heating systems. Due to these reasons, AdaHeat can be considered as a general framework where specific models can be inserted to give particular characteristics. Last but not least, AdaHeat adapts to the user preferences in balancing cost and discomfort as it relies on a single parametrization factor that is learned on-line.

In more detail, we extend the state-of-the-art as follows:

- We show how adaptive gray-box thermal modeling (i.e., adaptive modeling that relies on simplified physical equations—see Sec. 2) can be incorporated in IDHAs to capture the highly dynamic nature of domestic thermal characteristics. This is the first IDHA that incorporates adaptive gray-box thermal modeling.
- We propose a general algorithm for planning in the context of MPC, that optimally accounts for the occupancy probabilities and efficiently searches over the heating schedule space, utilizing dynamic programming.
- We evaluate our approach with data coming from a real house that employs underfloor heating (which constitutes a challenging testbed on the generality of our approach both in terms of thermal modeling and control) where we show the benefits of incorporating adaptive gray-box thermal modeling in IDHAs as well as the effectiveness of balancing heating cost and thermal discomfort based on a single parameter.

- We run a comparison over existing heating agent approaches and an improved approach that fully exploits the occupancy probabilities where we show that the latter leads to a more stable performance, in terms of Pareto efficiency, in various operational settings. In this context we also provide significant insights into the agents’ usability in various settings.

The rest of the paper is structured as follows: We begin with Sec. 2 with a review of related work. Then, in Sec. 3, we present our general adaptive IDHA, AdaHeat. In Sec. 4.4 we evaluate AdaHeat and provide a comprehensive characterization of IDHAs. Finally, Sec. 5 concludes.

## 2. RELATED WORK

As discussed above, two fundamental tasks in IDHAs are: i) the reliable thermal modeling of the house; and ii) the efficient control to optimize the heating process handling the occupancy uncertainty. Now, *adaptive* thermal modeling—where the thermal model varies through time to adapt to the dynamic thermal characteristics—has been shown to be resilient and effective in highly dynamic thermal settings such as those in houses (e.g., [20, 25, 12, 17, 23]). Based on the techniques used, adaptive modeling can be classified as either *black-box* or *gray-box* [25, 20]. Adaptive *black-box* approaches use statistical or machine learning techniques for thermal modeling with minimum need for prior physical knowledge of the system (e.g., [17, 23]). However, these approaches lack a physical interpretation and, hence, typically require a large amount of training data to demonstrate adequate performance (e.g., [20, 23]). Given this, such approaches are not considered in this work. On the other hand, adaptive *gray-box* thermal modeling, which we consider here, relies on simplified physical equations based on derived equivalent thermal parameters (ETPs) which are learned on-line and assumed to be time-varying (e.g., [20, 12]). Hence, such approaches do not suffer from the above drawbacks [20] and have already shown their potential in thermal processes control (e.g., [7, 5]). However, their incorporation in IDHAs has not been investigated yet.

Regarding control, a wide family of respective approaches that has proven very efficient and has been extensively used in IDHAs is that of MPC (e.g., [15, 28, 13, 31, 24]). This success of MPC is due to its ability to handle control problems where off-line computation of a control law is difficult or impossible (as is the case in IDHAs<sup>2</sup>) [4]. Moreover, the slow nature of thermal processes of buildings does not generally raise stability issues that are usually a concern with MPC control [30]. Due to these reasons, MPC is the most common approach in heating agents [10], and is also employed in this work. In more detail, MPC considers a wide family of control algorithms that share the following three criteria [4]: (i) they make explicit use of a model that describes the dynamics of the system under control in order to predict its future state; based on this model, (ii) they calculate a sequence of actions over a finite horizon according to the optimization objective—in this work we refer to this process as *planning*; and, finally, (iii) they apply the first control action of the calculated sequence, and repeat the

<sup>2</sup>We note here that although [29] calculates a control law off-line, it is impractical as it fails to consider real-time updates of the occupancy schedule and weather condition estimates.

procedure, shifting the planning horizon into the future—a property known as receding horizon.

A number of IDHAs have been proposed in the literature (e.g., [24, 28, 21, 31, 29]). In their pioneering work, [24] propose *Neurothermostat*, which employs a control method that fully exploits the occupancy probabilities and balances cost and discomfort in a single-objective optimization manner. However, the major drawback of this work is that it employs exhaustive search for planning which is extremely costly, limiting its applicability to simple proof-of-concept settings. In this context, *Neurothermostat* relies on a simple, fixed and, thus, impractical thermal model. Moreover, *Neurothermostat* employs a static empirical formula to express discomfort in monetary cost which is problematic as this equivalence varies among users and through time [28].

In contrast to the infinite horizon approach above, [21] propose *Smart Thermostat*, which divides heating control into two relatively independent tasks: (i) when to switch on heating; *preheating*, and (ii) when to switch it off; *heating stopping*. In this context, *Smart Thermostat* uses a simple, fixed thermal model based on ETPs estimated through historical average. However, *Smart Thermostat* employs a system-specific heuristic preheating approach that only searches over a sub-region of the heating schedule space. Moreover, heating stopping is reactive (heating is switched-off when a departure event is inferred) which is not Pareto optimal for heating systems that exhibit considerable thermal lags [11]. Most importantly though, tackling preheat and heating stopping independently is not effective in heating systems with considerable thermal lags (even if early stopping is considered), as the preheating policy can significantly affect the optimal stopping policy and vice-versa.

More recently, a particular heuristic approach is rising in popularity which deals with the probabilistic occupancy estimates in a thresholding manner (e.g. [28, 14, 13]). In particular, these estimates are assumed binary depending on their relation to a predefined threshold; any estimate above the threshold assumes occupancy, otherwise not. In more detail, [28] propose *PreHeat*, which plans based on a deterministic occupancy schedule, derived through the aforementioned approach. *PreHeat* works in two ways: (i) when the space is considered occupied it uses predefined set point temperatures; and if not (ii) it uses a lookahead window to check if an occupancy event is imminent so as to heat up the space for the minimum time required right before this event. In this context, *PreHeat* uses a simple fixed thermal model based on a single ETP, estimated as a historical average. That said, this method tackles preheating and heating stopping independently where the latter is reactive, thus facing the aforementioned limitations. Moreover, the preheating method is only appropriate for fixed-efficiency heating systems that do not exhibit any thermal lags, or energy cost variability over time. On top of the above, the trade-off between cost and discomfort is determined by the threshold choice which defines the deterministic occupancy schedule. As such, cost and discomfort are balanced based on a heuristic approach and no guarantees or intuition is given regarding the performance loss from a heating schedule planning that fully exploits the probabilistic occupancy estimates.

Another heating agent using thresholding is *SPOT+* [13], a non-domestic heating agent (for office buildings) that does deal with occupancy uncertainty and hence could also be employed in domestic settings. *SPOT+* tackles heating control

in an infinite horizon optimization manner but plans based on a threshold-based deterministic occupancy schedule. In this context, *SPOT+* uses a fixed thermal model for planning, estimated through least squares regression. By so doing, *SPOT+* balances discomfort and cost on two levels: (i) based on the threshold choice to derive with the deterministic occupancy schedule and, (ii) based on the weighting parameter used in the unifying formula. However, this scheme obscures how each of the balancing techniques affect the trade-off between cost and discomfort making parameter choice tricky. Moreover, this scheme also considers a heuristic approach and no intuition is given regarding the performance loss from optimal heating schedule planning. Lastly, although shortest path finding is mentioned for planning, the algorithmic choice is not reported and no appropriate algorithm is provided.

### 3. A GENERAL ADAPTIVE IDHA

In this work we propose a new general adaptive IDHA, *Ada-Heat* consisting of the following components: (i) the *thermal comfort model*, (ii) the *thermal model* of the building, (iii) the heating system *consumption model*, and (iv) the *controller*, that utilizes the aforementioned components. We now proceed to describe each component in detail.

#### 3.1 Thermal Comfort Model

In essence, thermal comfort is a complex response to several potentially interacting and less tangible factors (e.g., differences in mood, activity, biology, clothing, air temperature, humidity, and air speed) [8]. As such, based on different assumptions, a variety of metrics have been proposed to measure thermal discomfort [8]. In this work, for simplicity, we assume discomfort to depend only on the inside air temperature,  $T^{IN}$ ; and any discomfort experienced, at each instance that the house is occupied, to be the absolute deviation of  $T^{IN}$  from the user-provided, set-point temperature,  $T^{SP}$ . As such, assuming constant  $T^{IN}$  within a particular interval of length  $\delta$ , thermal discomfort is calculated as:

$$\text{Disc}(T^{SP}, T^{IN}, \delta) = |T^{SP} - T^{IN}| \delta \mathbb{1}_{\text{occupied}} \quad (1)$$

That said, more complex, and potentially self-adaptive, comfort models can be incorporated in our approach in a straightforward manner. Now, the occupancy schedule is essential for modeling and predicting thermal discomfort since the latter is experienced only when the house is occupied. However, the occupancy schedule is usually unknown in domestic settings and needs to be predicted. In this work, we employ the schedule-based occupancy prediction approach proposed by [28], due to its general low instrumentation needs and its particular efficiency (i.e. median predictive accuracies of  $\sim 80\%$ ) [28, 19]. In more detail, this approach predicts the occupancy schedule on-line and returns a vector of occupancy probabilities for every 15 min. over the predicting horizon.<sup>3</sup> As such, assuming a constant  $T^{IN}$  during an interval of length  $\delta$ , the expected thermal discomfort is:

$$\mathbb{E} [\text{Disc}(T^{SP}, T^{IN}, \delta)] = B |T^{SP} - T^{IN}| \delta \quad (2)$$

where  $B$  is the occupancy probability during the interval.

<sup>3</sup>We interpolate any estimates where necessary.

## 3.2 Thermal Model

In general, a thermal model predicts the thermal response of a building based on: (i) the current thermal state vector of the building,  $\mathbf{x}$ ; (ii) the vector of heating control actions to be executed,  $\mathbf{u}$ ; and (iii) the vector of information variables regarding exogenous stochastic processes that affect the thermal process (e.g, incident solar radiation, outside or adjacent buildings’ temperature),  $\mathbf{i}$ . As such, at time step  $t$ , any thermal model can be defined as  $\mathbf{x}_{t+1} = \mathcal{T}\mathcal{M}(\mathbf{x}_t, \mathbf{i}_t, \mathbf{u}_t)$ , where its parameters are the ETPs to be estimated which can be assumed to be either time-varying or fixed.

Now, in order to account for the highly dynamic domestic thermal characteristics (Sec. 2), we assume time-varying ETPs. Depending on the complexity of the thermal model used (linear or non-linear), different ETP identification methods can be used such as recursive least square with forgetting factor or (extended) Kalman filters.<sup>4</sup> We note, that although several methodologies exist for model selection (e.g., [26, 2]), identifying the most suitable model depends on the process to be modeled and the application requirements and, hence, it is typically undertaken by the designer. Nevertheless, our agent is able to handle both linear and non-linear models (as further discussed in Sec. 3.4). A specific instantiation for our case study system is provided in Sec. 4.3.

## 3.3 Consumption Model

Strictly speaking, the amount of energy consumed by space heaters is the energy provided to the space over the efficiency of the system. As such, we calculate the consumption cost over a time interval of length  $\delta$ , where its efficiency,  $C_{eff}$ , the energy price,  $P^{Buy}$ , and  $\mathbf{u}$ , remain constant, as:

$$\text{Cost}(\mathbf{u}, P^{Buy}, \delta) = \delta \frac{\text{Pwr}(\mathbf{u})}{C_{eff}} P^{Buy} \quad (3)$$

where,  $\text{Pwr}(\mathbf{u})$  stands for the energy provided to the space according to  $\mathbf{u}$  (and assumed to be independent of  $\mathbf{x}$  in terms of simplicity). Now, in contrast to other approaches where fixed formulas or multiple user-provided parameters are used to balance cost and discomfort (e.g., [24, 14]), in our approach this balancing is adaptive to the user preferences (see Sec. 3.4). As such,  $C_{eff}$ , as well as the energy provided for a particular  $\mathbf{u}$ , can be set to arbitrary values as long as any needed ratios are retained.<sup>5</sup> We note however that, for heat pumps,  $C_{eff}$  needs to be modeled as a function of the temperature difference between the heat source and the sink.

## 3.4 Control Approach

We now describe our control approach that utilizes the aforementioned components and provide our general planning algorithm. AdaHeat employs an adaptive, *certainty equivalent* MPC [3, 4] that works as follows: Every  $\delta$  amount of time, the controller executes the first action of the planned heating schedule. Then the thermal model is updated, new probabilistic occupancy estimates, and predictions of  $\mathbf{i}$  and, potentially,  $P^{Buy}$  are acquired; and the procedure is repeated shifting the planning horizon into the future. We now proceed to describe our planning approach.

<sup>4</sup>We note here that  $\mathbf{x}$  can be partially observable as well and estimating state variables along with the ETPs introduces non-linearity even in the case of a linear model [16].

<sup>5</sup>For instance, for a radiative heater which can operate with either 1 or 2 identical elements, the respective energy provision values should correspond to the ratio 0:1:2.

## Planning (Objective Formalization).

Within the above defined context, planning considers the task of balancing discomfort and cost over the planning horizon; that is a finite bi-objective optimization problem. In order to tackle the respective complexity we combine the two objectives via the well-known *weighted sum* [22] which is a *sufficient* but not necessary condition for Pareto optimality [22].<sup>6</sup> Although other formalizations exist that are both necessary and sufficient conditions [22], we employ the weighted sum due to its simplicity and good observed performance. Moreover, by doing so, our agent is able to adapt to the user’s preferences, through a simple Boolean feedback procedure. In particular, the user can simply progressively adjust the weighting factor, in real time, by a constant value until his/her preferences are met.

More formally, we plan for the MPC horizon, of length  $\Delta$ , by breaking it down into a set of non-overlapping intervals of length  $\delta$ . As such, ensuring that  $\Delta$  is an integer multiple of  $\delta$ , it corresponds to a set of intervals, noted  $H$ , where  $|H| = \Delta/\delta$ . During each interval,  $\tau$ , all environmental conditions are assumed constant. Hence, the optimization objective is to find the sequence of actions,  $\mathbf{u}_\tau$ , that minimizes the expected unifying cost;  $\bar{J}$ , over the planning horizon:

$$\begin{aligned} & \underset{\mathbf{u}_1, \dots, \mathbf{u}_{|H|}}{\text{minimize}} && \bar{J}(\cdot) = \sum_{\tau=1}^{|H|} \lambda \mathbb{E}[\text{Disc}(\cdot)] + (1 - \lambda) \text{Cost}(\cdot) \\ & \text{subject to} && \mathbf{u}_1, \dots, \mathbf{u}_{|H|} \in U \end{aligned}$$

Here,  $\mathbb{E}[\text{Disc}(\cdot)]$  and  $\text{Cost}(\cdot)$  return the expected discomfort and heating cost during interval  $\tau$ , based on Eq. 2 and 3 respectively,  $\lambda$  is the weighting factor, and  $U$  is the set of all feasible  $\mathbf{u}$ .<sup>7</sup> In general,  $\lambda \in (0, 1)$ , to ensure strict Pareto optimality [22]—as for the two limits only one objective is considered. For eliminating cost with the minimum discomfort there is a unique trivial solution of no heating, while for the reverse problem a  $\lambda$  very close to 1 can be used. Note that normalized values for cost and discomfort can be used. We now proceed to describe our planning algorithm.

## Planning (Optimization Approach).

The slow nature of the thermal process of buildings enables us to tackle the optimization problem in a dynamic programming manner (as the real-time computation constraints are not typically strict). By doing so, our IDHA is general enough to handle both linear and non-linear models and objectives. In particular, we reduce planning into finding the shortest path in a directed acyclic graph (DAG) and provide a planning algorithm that exploits the property of topological ordering of a DAG through depth first search (DFS) to find the shortest path in linear time [6].<sup>8</sup>

In more detail, each node,  $n$ , of the DAG,  $G$ , corresponds to a distinct tuple that contains all the necessary information to predict the next state,  $n'$ , according to  $\mathbf{u}$  based on  $\mathcal{T}\mathcal{M}(\cdot)$ . We note though that  $\mathbf{i}$  at each instance can be inferred by  $\tau$  and, hence, the  $n$  tuple will just be  $\langle \tau_n, \mathbf{x}_n \rangle$ . Now, each of the edges,  $e$ , corresponds to a tuple that contains the initial

<sup>6</sup>Unless the Pareto optimal hyper-surface is convex [22].

<sup>7</sup>Note that the absence of hard output variable constraints ensures that we will not face any feasibility issues [4].

<sup>8</sup>Although dynamic programming requires a discretization, due to the limited predicting ability of any thermal model, a discretization of  $\mathbf{x}$  comes naturally (in contrast to the claim in [24]) while  $\tau$  is already discrete.

---

**Algorithm 1** AdaHeat Planning Algorithm

---

```
1: procedure HEATINGPLANNING( $G, n$ )
2:   for every  $u \in U$  do
3:      $Cost \leftarrow Cost(u, P^{Buy}[\tau_n], \delta)$ 
4:      $\overline{Disc} \leftarrow B[\tau_n] \text{ Disc}(T^{SP}, T_n^{IN}, \delta)$ 
5:      $n' \leftarrow \mathcal{TM}(n, u)$ 
6:      $e \leftarrow n \mapsto n', Cost, \overline{Disc} >$ 
7:     add  $e$  to  $E_G$ 
8:     if  $n' \notin V_G$  then
9:       add  $n'$  to  $V_G$ 
10:      if  $\tau_{n'} < |H|$  then
11:         $G \leftarrow \text{HEATINGPLANNING}(G, n')$ 
12:      else
13:         $\text{MinJ}\{n'\} \leftarrow 0$ 
14:         $\text{Tmp} \leftarrow \text{MinJ}\{n'\} + \lambda \overline{Disc} + (1 - \lambda) Cost$ 
15:        if  $\text{MinJ}\{n\} = \text{NaN}$  or  $\text{MinJ}\{n\} > \text{Tmp}$  then
16:           $\text{MinJ}\{n\} \leftarrow \text{Tmp}$ 
17:           $\text{BestAction}\{n\} \leftarrow u$ 
18:   return  $G$ 
```

---

and successor node, and two weights corresponding to the heating cost and expected discomfort during interval  $\tau$ .

Algorithm 1 illustrates our planning algorithm. In particular, we extend the DFS recursion with constant time expressions, thus the time complexity is retained at  $O(|V_G| + |E_G|)$  where  $V_G$  and  $E_G$  stand for the set of edges and vertices of  $G$  respectively. Specifically, the algorithm creates the DAG in a pre-order [6] manner and populates  $\text{MinJ}\{n\}$  and  $\text{BestAction}\{n\}$  with the minimum additional expected unifying cost,  $J$  and the best action for each node respectively. As such when the algorithm terminates  $\text{BestAction}\{n\}$  holds the optimal heating actions for each node.<sup>9</sup>

## 4. EVALUATION

In this section we provide a thorough evaluation of AdaHeat and a comprehensive characterization of state-of-the-art IDHAs; we first describe the case study of our evaluation and how we collected the necessary data; then, we describe the specific instantiation of our IDHA for the case study system; subsequently, we discuss our evaluation set-up and the instantiations of the various benchmark agents; and, then, we report the evaluation results.

### 4.1 Case Study and Data Collection

For our evaluation case study, we consider the living room of a family house in Cambridge, UK. The house has both radiators and underfloor heating (UFH) and is equipped with custom hardware for heating control and data collection (see [28]). We chose the living room as: (i) it is often in use and (ii) its thermal dynamics are particularly challenging due to its physical properties and household activity. In particular, it has two doors and three windows and it is equipped with a UFH system and an, occasionally used, auxiliary fan heater. As such, the heating in adjacent rooms, the weather conditions and the occupant activity have a substantial effect on its thermal dynamics. Moreover, UFH involves multiple heat transfer processes introducing considerable thermal lags. Taken together, these factors make this room a challenging testbed on the generality and efficiency of our approach both in terms of thermal modeling and control.

<sup>9</sup>The arguments at the initial call of the recursion consider an empty graph and the root node.

For the purpose of our research, we collected  $T^{IN}$  readings, and occupancy events from November 2011 to March 2012 (150 days) via the custom hardware (as discussed above). For the outside temperature,  $T^o$ , we use the publicly available dataset from the Cambridge Computer Laboratory.<sup>10</sup> Finally, for solar radiation estimates,  $G^s$ , we use the dataset from the EU Joint Research Commission.<sup>11</sup>

### 4.2 Instantiating AdaHeat

We now detail the instantiation of AdaHeat for our case study system. Starting with *comfort modeling*, the case study set-point temperature and, hence  $T^{SP}$ , is 22°C.

Regarding *thermal modeling*, we identify the most suitable model by starting with the simplest feasible model and iteratively refining it into a more complex one. By doing so, we derive a model where the transfer of heat from the source to the indoor air is assumed to take place via an intermediate thermal mass, and the transfer to the outside via the house envelope. Moreover, our model accounts for the effects of solar radiation on the indoor air and house envelope temperature. In more detail, the derived model is [1]:

$$\begin{aligned} T_{t+1}^{FL} &= T_t^{FL} + r^h a + \phi_a (T_t^{IN} - T_t^{FL}) \\ T_{t+1}^{IN} &= T_t^{IN} + r_a^s G^s + \phi_a (T_t^{FL} - T_t^{IN}) + \phi_b (T_t^{EN} - T_t^{IN}) \\ T_{t+1}^{EN} &= T_t^{EN} + r_e^s G^s + \phi_b (T_t^{IN} - T_t^{EN}) + \phi_c (T^o - T_t^{EN}) \end{aligned}$$

where  $T^{FL}$  and  $T^{EN}$  stand for the floor-mass and envelope temperature respectively, and, along with  $T^{IN}$ , consider  $\mathbf{x}$ . Furthermore,  $G^s$  and  $T^o$  consider  $\mathbf{i}$ . In addition,  $\phi_a$ ,  $\phi_b$ , and  $\phi_c$  stand for leakage rates<sup>12</sup>, and  $r^h$ ,  $r_a^s$  and  $r_e^s$  are additional coefficients that capture the effect of the heating output on  $T^{IN}$ , and the effect of  $G^s$  on  $T^{IN}$  and  $T^{EN}$ , respectively. These coefficients along with the leakage rates consider the time-varying ETPs. Finally,  $a \in \{1, 0\}$  (on/off) is the heating action and trivially considers  $\mathbf{u}$ . That said,  $T^{FL}$  and  $T^{EN}$  are hidden thermal state variables that need to be estimated along with the ETPs. To this end, as is common practice, we use an extended Kalman filter (EKF) for the joint estimation of state and parameter variables [16], and evaluate our approach over the 150-days dataset to achieve the 95th percentile of the absolute prediction error to be 0.95°C and 1.37°C for 2 and 4 hours predictions, respectively.

Regarding the *consumption model*, we have appropriately set  $C_{eff} = 1$  and  $\text{Pwr}(a) = a$  (i.e.,  $\text{Pwr}(0) = 0$  and  $\text{Pwr}(1) = 1$ ). Finally, the planning horizon and the planning interval length of AdaHeat *controller* were set to 1 hour ahead and 5 minutes, respectively (i.e.,  $\delta = 5$  min,  $|H| = 12$ ), as those MPC design characteristics have been found to be adequate for efficient control of the case study system (after experimenting with various design characteristics).

### 4.3 Experimental Setup

In this work, we evaluate AdaHeat, *with* and *without* adaptive thermal modeling. We do so, to identify the benefits of such modeling in our IDHA. Moreover, we compare against the well-known SPOT+ and PreHeat which, essentially, employ MPC along with heuristic planning (Sec. 2). As such,

<sup>10</sup>[www.cl.cam.ac.uk/research/dtg/weather](http://www.cl.cam.ac.uk/research/dtg/weather)

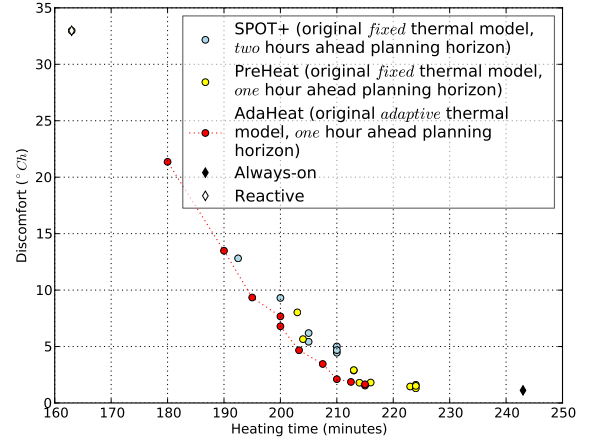
<sup>11</sup>[re.jrc.ec.europa.eu/pvgis/apps4/pvest.php](http://re.jrc.ec.europa.eu/pvgis/apps4/pvest.php)

<sup>12</sup>As far as the common RC-network representation is considered [20], the notion of leakage rates can be interpreted as the cumulative representation of thermal capacitance,  $C_{th}$ , and thermal resistance,  $R_{th}$  (i.e.,  $\phi = \frac{1}{C_{th}R_{th}}$ ).

our evaluation can provide significant insights about the trade-off between heuristic and optimal planning, in the context of MPC. Now, although these heating agents employ simple fixed thermal modeling, we also evaluate them with a more advanced *fixed* model (that captures the thermal lags of the case study system) and with our *adaptive* model. In addition, we use the same occupancy prediction algorithm (i.e., [28]); cost and discomfort metrics (i.e., Eq. 3 and 1 respectively); and planning horizon and interval length (i.e.,  $\delta = 5$  min and  $|H| = 12$ ) for all agents. We do so: (i) to identify the benefits of adaptive modeling in various IDHAs; and (ii) to compare various IDHAs without being affected by any model and design differences. Moreover, we evaluate all IDHAs *with* and *without* considering variable energy cost in order to characterize them in different settings. That said, our case study system with energy cost variability is a worst case scenario system and its efficient control can confirm (or disprove) the intended generality of AdaHeat. For completeness, we also evaluate the performance of three simple heating strategies: (i) Always-on, which retains  $T^{IN}$  at  $T^{SP}$  throughout the whole day, (ii) Never/Always-off, in which heating is always off, and (iii) Reactive, in which heating responds to occupancy (this is equivalent to a strategy where heating is manually switched on and off, when the occupants leave and return to the house, respectively). In more detail, the aims of this evaluation are: (i) to identify the benefits of incorporating adaptive gray-box thermal modeling in different IDHA approaches, (ii) to identify the trade-off between heuristic planning and a planning approach that fully exploits the probabilistic occupancy estimates, in the context of MPC (without being affected by any modeling and design differences of the IDHAs considered), and (iii) to provide a comprehensive comparison of different IDHAs in different operational settings (also without being affected by any modeling and design differences).

In more detail, we evaluate all IDHAs for a typical *winter* day (of February 2011), ensuring (via an iterative procedure) that the initial and final thermal state,  $\mathbf{x}$ , at the beginning and at the end of the day respectively, are the same for all experiments. As such, our evaluation results consider *long-term average performance evaluation*, assuming that the same day repeats over time (i.e., same occupancy schedule, environmental conditions and predictions). We followed this procedure to provide *long-term performance estimates* for *various* IDHA parameter settings within feasible computational time. In particular, by doing so, we were able to evaluate all IDHAs for a wide range of parameters and identify their performance in meeting the user preferences. In more detail, we evaluate SPOT+ for all combinations of a weighting factor within (0,1) with a step of 0.01, and a threshold value within (0,1) with step 0.1.<sup>13</sup> In addition, we evaluate PreHeat and AdaHeat for the same threshold and weighting factor range, respectively.

Now, we chose a week-day in winter due to the heating needs of the particular season and to avoid any week-end day peculiar features.<sup>14</sup> To this end, we used the collected data for the ground truth of the occupancy schedule (and derived respective occupancy predictions for this day based



**Figure 1: Initial evaluation results**<sup>17</sup> (Always-off yields  $\sim 154^\circ\text{C/h}$  discomfort and, appropriately, no cost)

on historical data according to [28]—see Sec. 3.1) and the weather conditions (see Sec. 4.1).<sup>15</sup> Furthermore, in order to model our thermal model inaccuracies, we simulated the underlying thermal process by sampling  $\mathbf{x}$ , at each instance, from the respective EKF derived distributions. As such, the thermal model is not completely accurate with respect to our simulation, making our experiments more realistic.

As outlined above, we evaluate SPOT+ and PreHeat also with their original fixed thermal models. Thus, as proposed in the respective publications, we estimated SPOT+ model via least squares regression and PreHeat’s heat-rate as a historical average. In particular, we estimated both models based on the two first months of the 150-days dataset—thereafter, the estimated ETPs are fixed. Now, in order to evaluate the IDHAs with a more advanced fixed model, we used the ETPs of our adaptive model as derived exactly 30 days before the evaluation day (as such, the last model “calibration” is done, approximately, one month ago).<sup>16</sup> Finally, we note that in our evaluation, discomfort and cost are estimated with a numerical evaluation of *1-min* interval.

## 4.4 Evaluation Results

We note that the simple models of SPOT+ and PreHeat are not able to capture the case study UFH system thermal lags. However, as further discussed below, PreHeat is not very sensitive to the accuracy of the thermal model used due its simple heating control strategy. On the other hand though, SPOT+ is not able to execute any heating schedule other than Always-off when a planning horizon of one hour is used. Hence, only for this experiment SPOT+ horizon is set to *two* hours (in contrast to AdaHeat and PreHeat where *one* hour is used). Fig.1 illustrates the evaluation results.

From this we can see that AdaHeat has a better performance, in terms of Pareto efficiency, compared to both SPOT+ (with two hours ahead planning horizon) and PreHeat (while SPOT+ and PreHeat have a comparable ef-


<sup>13</sup>We note here that the SPOT+ objective has been normalized to work with a weighting parameter in the range (0,1) *without any performance loss* to reinforce our comparison.

<sup>14</sup>For instance if the house is unoccupied during a week-end day, due to a trip, there will be zero potential savings.

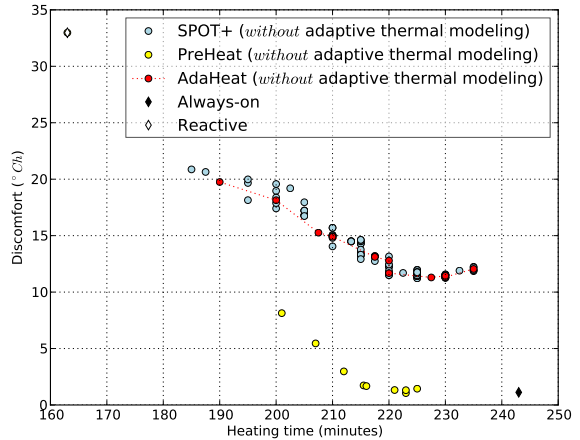
<sup>15</sup>We linear interpolate whenever needed.

<sup>16</sup>Although this simple technique is used to “approximate” a fixed thermal model, estimation techniques for fixed ETPs can potentially demonstrate higher accuracy [18].

<sup>17</sup>Points closer to the origin indicate higher Pareto efficiency.



(a) *with* adaptive thermal modeling



(b) *without* adaptive thermal modeling

**Figure 2: Comprehensive evaluation results**<sup>17</sup> (Always-off yields  $\sim 154^\circ\text{Ch}$  discomfort and no cost)

efficiency). In particular, the balancing points captured by AdaHeat fall closer to the origin and consider a wider and more evenly distributed set. However, this experiment is not very informative on whether this is due to the differences in thermal modeling or in planning. It is worth noting though, that none of the agents are dominated by the simple strategies (i.e., Always-on, Always-off or Reactive) and can improve heating efficiency compared to these strategies.

Given these initial observations, we proceed with a more comprehensive evaluation of the agents. In particular, we first evaluate them with our thermal modeling approach, both fixed and adaptive, without considering energy cost variability. As expected, adaptive thermal modeling significantly improves IDHA efficiency. In particular, as seen in Fig. 2, both AdaHeat and SPOT+ highly depend on the thermal model accuracy and their performance improves significantly when adaptive modeling is considered, especially when low discomfort values are intended. In particular, the solutions captured with adaptive modeling fall closer to the origin compared to fixed modeling solutions. On the other hand, PreHeat is less sensitive to the thermal model accuracy due to its simple control strategy. However, this simple strategy deteriorates in terms of flexibility and efficiency as discussed below. In general though, none of the systems' solutions are dominated by the simple heating strategies even when fixed thermal modeling is considered.

Now, as far as all IDHAs are considered with adaptive thermal modeling, the results shown in Fig. 2(a) suggest that SPOT+ and AdaHeat have comparable Pareto efficiency while PreHeat demonstrates a slightly worse efficiency. This is due to its simple control strategy which is not able to capture heating systems with considerable thermal lags, such as the UFH system considered, in a maximally efficient manner. Moreover, SPOT+ demonstrates a less stable performance, in terms of Pareto efficiency compared to AdaHeat, i.e., the solutions captured by SPOT+ are sometimes dominated by AdaHeat and vice-versa. In further investigation, SPOT+ has been observed to occasionally plan a clearly suboptimal heating schedule.<sup>18</sup> However, the suboptimal planning

of SPOT+ occasionally leads to higher or lower Pareto efficiency, as the MPC is not an optimal control approach [4].

Now, matching the time-varying occupant preferences in balancing discomfort and cost is crucial in the context of IDHAs (as discussed in Sec. 2). To this end, SPOT+ relies on two user-provided parameters, i.e., the weighting factor and the threshold over the probabilistic occupancy estimates (Sec. 2). However, in general, mathematical relationships between heating cost and quantifications of thermal discomfort are hard to comprehend for the users. As such, the usability of SPOT+ in domestic settings is questionable due to the complicated relationship between the threshold and the weighting parameter (see Fig. 3(a)). In more detail, many SPOT+ solutions (for different weighting and threshold parameters) are dominated by other solutions that SPOT+ captures with different parameter choices. However, the exact performance of SPOT+ cannot be known in advance and we are not able to find any algorithm to appropriately populate the weight and the threshold parameter that can demonstrate a monotonic relationship with either the discomfort or the cost. For instance, one such algorithm could be to increase weight and threshold iteratively, starting from a particular weight for each threshold choice. This fact makes the parameter choice tricky as the user cannot know what to expect from different parameter value combinations.

On the other hand, both AdaHeat and PreHeat rely on only a single parameter for balancing heating cost and thermal discomfort. Moreover, the adjustable parameter of both AdaHeat and PreHeat demonstrates a monotonic correlation to thermal discomfort (and to heating cost for AdaHeat), when adaptive thermal modeling is considered, as seen in Figures 3(b) and 3(c). This fact enables the adjustment of these variables through a simple, real-time, adaptive procedure, based on a single Boolean feedback from the user, as discussed in Section 3.4. As discussed above though, when adaptive thermal modeling is considered (Fig. 2(a)), PreHeat illustrates a slightly lower Pareto efficiency than

of SPOT+'s planning objective (see [13]) is not a sufficient condition for Pareto optimality over cost and expected discomfort. However, we cannot conclude, whether it is a necessary condition just from these observations.

<sup>18</sup>This fact suggests that the non-closed form formalization




Figure 3: Balancing heating cost and thermal discomfort<sup>19</sup>

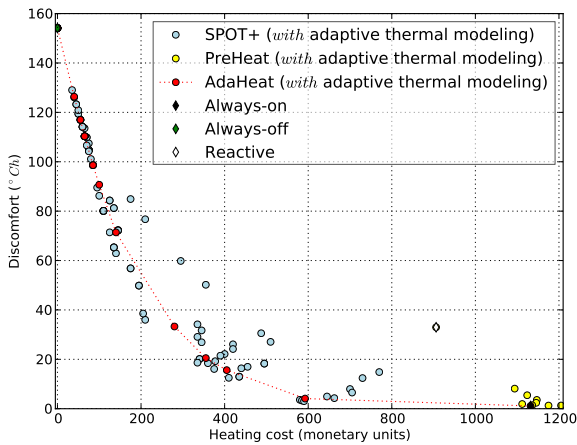


Figure 4: Evaluation results with variable price<sup>17</sup>

SPOT+ and AdaHeat. Moreover, in general, PreHeat is not able to capture a wide range of balancing points between cost and discomfort that allows a variety of user preference schemes to be captured—in contrast to AdaHeat. In particular, PreHeat operates on only a small region in balancing cost and discomfort which is not sufficient for appropriate heating control in domestic settings. As such, the occupants need to also adjust the origin of the discomfort metric (i.e., the set-point temperature), along with the threshold parameter, in order to meet their preferences. Thus AdaHeat is the only agent that works sufficiently based on a single weighting parameter that can be learned on-line.

Lastly, the simple heating control strategy of PreHeat (i.e., heating for the minimum time required right before an occupancy event) does not allow this agent approach to efficiently work in conjunction with heating systems that exhibit a variability of heating energy cost over time, time-varying overall efficiency or considerable thermal lags (as illustrated above). To further illustrate this we have conducted an additional experiment where arbitrarily variable energy prices have been assumed through the day. In particular, the energy prices have been designed to change every 5

minutes with their value being sampled from a uniform distribution within the range [1,10]. As can be seen in Fig. 4, PreHeat’s performance deteriorates significantly in this settings (both in terms of Pareto efficiency, and distribution and range of balancing points that it captures). Specifically, certain PreHeat solutions are dominated even by the Always-on strategy. Moreover, SPOT+ demonstrates significant variability over its performance for different parameter choices in this setting, as it captures many self-dominated solutions. In contrast, AdaHeat is generally stable in terms of Pareto efficiency, and generally smooth in terms of the distribution and range of solutions captured (Fig. 4).

## 5. CONCLUSIONS

In this work we propose a new general IDHA (framework), AdaHeat, that balances cost and discomfort in an infinite horizon optimization manner, learns an adaptive thermal model on-line and does planning to fully exploit the occupancy probabilities. AdaHeat adapts to the user preferences in balancing cost and discomfort as it relies on only one parametrization factor. We showed the effectiveness of our approach in different settings against two state-of-the-art IDHAs. In particular, we showed how adaptive thermal modeling can significantly improve the efficiency of IDHA, especially when advanced heating strategies are considered (i.e., SPOT+, AdaHeat). Moreover, we showed that a single parameter, in balancing heating cost and thermal discomfort, is sufficient for efficient IDHA performance, and ensures the applicability of the IDHAs in domestic settings with variable user preferences. In addition, we showed that optimal exploitation of the occupancy probabilistic estimates, within the context of MPC, is feasible in practice and leads to a more stable performance, in terms of Pareto efficiency, in various operational settings. Lastly, in this work, we also ran a comparison over existing IDHAs and provided significant insights about their performance in terms of Pareto efficiency and usability criteria. Regarding future work we aim to extended our work to control heating in systems with renewable energy resources and uncertain energy prices since such systems are expected to be the norm within the next generation electricity grid (i.e., the smart grid [27]).

**Acknowledgements** We thank Lampros Stavrogianis (University of Southampton) for comments and fruitful discussions. We also acknowledge funding from the UK Research Council for the ORCHID project, grant EP/I011587/1.

<sup>19</sup> Fig. 3(a) reveals the complicated, non-monotonic and, hence, impractical nature of balancing discomfort and cost based on two parameters (in contrast to 3(b) and 3(c)).



## REFERENCES

- [1] M. Alam, A. A. Panagopoulos, A. Rogers, N. R. Jennings, and J. Scott. Applying extended kalman filters to adaptive thermal modelling in homes: Poster abstract. In *Proceedings of the 1st ACM Conference on Embedded Systems for Energy-Efficient Buildings*, BuildSys '14, pages 214–215, NY, USA, 2014. ACM.
- [2] P. Bacher and H. Madsen. Identifying suitable models for the heat dynamics of buildings. *Energy and Buildings*, 43(7):1511–1522, 2011.
- [3] D. P. Bertsekas. Dynamic programming and suboptimal control: A survey from ADP to MPC. *European Journal of Control*, 11(4):310–334, 2005.
- [4] E. F. Camacho and C. B. Alba. *Model predictive control*. Springer, 2013.
- [5] T. Chen. Application of adaptive predictive control to a floor heating system with a large thermal lag. *Energy and Buildings*, 34(1):45 – 51, 2002.
- [6] T. H. Cormen, C. E. Leiserson, R. L. Rivest, C. Stein, et al. *Introduction to algorithms*. MIT press, 2001.
- [7] J. B. Cunha, C. Couto, and A. Ruano. Real-time parameter estimation of dynamic temperature models for greenhouse environmental control. *Control Engineering Practice*, 5(10):1473 – 1481, 1997.
- [8] N. Djongyang, R. Tchinda, and D. Njomo. Thermal comfort: A review paper. *Renewable and Sustainable Energy Reviews*, 14(9):2626–2640, 2010.
- [9] B. Dong, K. P. Lam, and C. Neuman. Integrated building control based on occupant behavior pattern detection and local weather forecasting. In *Twelfth International IBPSA Conference. Sydney: IBPSA Australia*, pages 14–17, 2011.
- [10] A. I. Dounis and C. Caraiscos. Advanced control systems engineering for energy and comfort management in a building environment—a review. *Renewable and Sustainable Energy Reviews*, 13(6):1246–1261, 2009.
- [11] C. Ellis, J. Scott, M. Hazas, and J. Krumm. Earlyoff: Using house cooling rates to save energy. In *Proceedings of the Fourth ACM Workshop on Embedded Sensing Systems for Energy-Efficiency in Buildings*, pages 39–41. ACM, 2012.
- [12] S. Fux, A. Ashouri, M. Benz, and L. Guzzella. EKF based self-adaptive thermal model for a passive house. *Energy and Buildings*, 68, Part C(0):811 – 817, 2014.
- [13] P. X. Gao and S. Keshav. Optimal personal comfort management using SPOT+. In *Proceedings of the 5th ACM Workshop on Embedded Systems For Energy-Efficient Buildings*, pages 1–8. ACM, 2013.
- [14] P. X. Gao and S. Keshav. SPOT: A smart personalized office thermal control system. In *Proceedings of the Fourth International Conference on Future Energy Systems*, pages 237–246. ACM, 2013.
- [15] R. Halvgaard, N. K. Poulsen, H. Madsen, and J. B. Jorgensen. Economic model predictive control for building climate control in a smart grid. In *Innovative Smart Grid Technologies (ISGT), 2012 IEEE PES*, pages 1–6. IEEE, 2012.
- [16] S. S. Haykin. *Kalman filtering and neural networks*. John Wiley & Sons, Inc., New York, NY, USA, 2001.
- [17] H. Huang, L. Chen, M. Mohammadzaheri, E. Hu, and M. Chen. Multi-zone temperature prediction in a commercial building using artificial neural network model. In *10th IEEE Intern. Conference on Control and Automation (ICCA)*, pages 1896–1901, 2013.
- [18] K. Keesman. *System Identification: An Introduction*. Advanced Textbooks in Control and Signal Processing. Springer, 2011.
- [19] W. Kleiminger, F. Matterna, and S. Santinib. Predicting household occupancy for smart heating control: A comparative performance analysis of STOA approaches. Technical report, ETH Zurich, 2013.
- [20] X. Li and J. Wen. Review of building energy modeling for control and operation. *Renewable and Sustainable Energy Reviews*, 37(0):517 – 537, 2014.
- [21] J. Lu, T. Sookoor, V. Srinivasan, G. Gao, B. Holben, J. Stankovic, E. Field, and K. Whitehouse. The smart thermostat: Using occupancy sensors to save energy in homes. In *Proc. of the 8th ACM Conf. on Embedded Networked Sensor Systems*, pages 211–24. ACM, 2010.
- [22] R. T. Marler and J. S. Arora. Survey of multiobjective optimization methods for engineering. *Structural and multidisciplinary optimization*, 26(6):369–395, 2004.
- [23] N. Morel, M. Bauer, M. El-Khoury, and J. Krauss. Neurobat, a predictive and adaptive heating control system using artificial neural networks. *International Journal of Solar Energy*, 21(2-3):161–201, 2001.
- [24] M. C. Mozer, L. Vidmar, R. H. Dodier, et al. The neurothermostat: Predictive optimal control of residential heating systems. *Advances in Neural Information Processing Systems*, pages 953–959, 1997.
- [25] S. Privara, J. Cigler, Z. Vana, F. Oldewurtel, C. Sagerschnig, and E. Zacekova. Building modeling as a crucial part for building predictive control. *Energy and Buildings*, 56(0):8 – 22, 2013.
- [26] S. Privara, Z. Vána, E. Žáčková, and J. Cigler. Building modeling: Selection of the most appropriate model for predictive control. *Energy and Buildings*, 55:341–350, 2012.
- [27] S. D. Ramchurn, P. Vytelingum, A. Rogers, and N. R. Jennings. Putting the smarts' into the smart grid: A grand challenge for artificial intelligence. *Communications of the ACM*, 55(4):86–97, 2012.
- [28] J. Scott, A. Bernheim Brush, J. Krumm, B. Meyers, M. Hazas, S. Hodges, and N. Villar. Preheat: controlling home heating using occupancy prediction. In *Proceedings of the 13th International Conference on Ubiquitous Computing*, pages 281–290. ACM, 2011.
- [29] M. Shann and S. Seuken. Adaptive home heating under weather and price uncertainty using GPS and MDPS. In *Proc. of the 2014 International Conf. on Autonomous Agents and Multi-agent Systems (AAMAS)*, pages 821–828. Intern. Foundation for Autonomous Agents and Multiagent Systems, 2014.
- [30] J. Široký, F. Oldewurtel, J. Cigler, and S. Privara. Experimental analysis of model predictive control for an energy efficient building heating system. *Applied Energy*, 88(9):3079–3087, 2011.
- [31] D. Urieli and P. Stone. A learning agent for heat-pump thermostat control. In *Proc. of the 2013 Int. Conf. on Autonomous Agents and Multi-agent Systems (AAMAS)*, pages 1093–100. Int. Found. for Autonomous Agents and Multiagent Systems, 2013.