Adaptive Home Heating Control Through Gaussian Process Prediction and Mathematical Programming

Alex Rogers, Sasan Maleki, Siddhartha Ghosh, Nicholas R. Jennings
School of Electronics and Computer Science
University of Southampton, SO17 1BJ, UK
{acr,sm9g09,sg2,nrj}@ecs.soton.ac.uk

ABSTRACT

In this paper, we address the challenge of adaptively controlling a home heating system in order to minimise cost and carbon emissions within a smart grid. Our home energy management agent learns the thermal properties of the home, and uses Gaussian processes to predict the environmental parameters over the next 24 hours, allowing it to provide real-time feedback to householders concerning the cost and carbon emissions of their heating preferences. Furthermore, we show how it can then use a mixed-integer quadratic program, or a computationally efficient greedy heuristic, to adapt to real-time cost and carbon intensity signals, adjusting the timing of heater use in order to satisfy preferences for comfort whilst minimising cost and carbon emissions. We evaluate our approach using weather and electricity grid data from January 2010 for the UK, and show our approach can predict the total cost and carbon emissions over a day to within 9%, and show that over the month it reduces cost and carbon emissions by 15%, and 9%, respectively, compared to using a conventional thermostat.

Categories and Subject Descriptors

I.2.11 [**Computing Methodologies**]: Distributed Artificial Intelligence

General Terms

Algorithms, Energy, Control, Learning

Keywords

Smart grid, machine learning, Gaussian process, mixed-integer optimisation, energy feedback, carbon emissions

1. INTRODUCTION

The creation of a smart electricity grid represents one of the greatest engineering challenges of this century, as countries face dwindling non-renewable energy sources and work to minimise the adverse effects of climate change due to carbon emissions [1, 2]. Key components of this vision include ambitious targets for renewable generation, the roll-out of smart meters to domestic consumers in order to facilitate real-time pricing of electricity, and a shift toward the electrification of heating through the use of air and ground source heat pumps.[1]

However, these developments present a potential challenge for householders, since the increased use of intermittent renewable generation means that both the price of electricity, and also the carbon intensity of the electricity (quoted in terms of gCO_2/kWh and signifying the amount of carbon dioxide emitted when one unit of electricity is consumed), will vary in real-time. These real-time signals will likely be passed to consumers, through their smart meters, who will be expected to respond rationally in order to reduce demand for expensive and carbon intensive electricity at peak times, and to make better use of low carbon renewable energy when it is available. However, the links between heater system control settings and energy consumption is already poorly understood by consumers [4], and these changes are likely to compound this issue. Thus, it is essential that future home heating systems are able to provide real-time feedback to householders concerning the implications (in terms of both cost and carbon emissions) of their heating preferences. Going further, these systems should adapt to these real-time signals, adjusting the timing of heater use in order to satisfy the home owners' preferences for comfort while also minimising cost and carbon emissions.

Now, the idea of individual homes responding to real-time signals from an electricity grid is not a new one, and indeed, was first proposed by [10], who discussed the scheduling of loads, and the prediction of both demand in the home, and local weather conditions. However, their work dealt mainly with predicting the overall system behaviour, and used closed form solutions that required approximating and homogenising the behaviours of the individual actors within the systems. As such, it does not actually provide a solution that can be implemented within any individual home. More recently, a number of researchers have revisited these ideas, and proposed solutions that address parts of the challenge described above. For example, [7] use artificial neural networks to predict the heat demand of individual homes using several external factors, but make no attempt to then optimise energy use. Conversely, [11] use mixed integer programming to optimise energy storage within a home which receives real-time price signals, but they consider electrical loads, rather than heating, and hence they do not need to consider the thermal characteristics of the home, nor the external factors that affect it.

Thus, to address these shortcomings, in this paper we develop a home energy management agent, or smart controller,

[1]The UK government has committed to reducing carbon emissions by 80% by 2050, and aims to install smart meters to all 26M UK homes by 2020 [6].

that works on behalf of the householder, and is able to learn the thermal characteristics of the home, and predict the additional factors that affect the cost and carbon emissions of heating use. It is then able to provide householders with the required real-time feedback, and also to autonomously control the heating system on their behalf. In more detail, using internal and external temperature sensors, and by monitoring the activity of the home's heating system, the smart controller finds model parameters that describe the heat output of the heating system and the thermal leakage rate of the home. In addition, using a Gaussian process model that exploits periodicities and correlations between time series, the smart controller predicts the local external temperature over the next 24 hours by combining local measurements from a sensor, with predictions from an online weather forecast. In doing so, it creates a *site-specific* forecast for the next 24 hours. Similarly, using available predictions of future demand, the smart controller predicts the carbon intensity of the electricity supplied to the home. Using these factors it can predict the consequences, in terms of cost and carbon, of any heater control setting and provide this information to the home owner through a graphical user interface. Going further, the smart controller is then able to fully optimise the use of heating (using either a mixed integer solver or a computationally efficient greedy heuristic). In doing so, it seeks to provide the same level of comfort as a standard thermostat operating at the same set-point temperature (evaluated using a comfort model based on the ANSI/ASHRAE Standard 55-2010) whilst also minimising either cost or carbon emissions.

Thus, in more detail, in this paper we make the following contributions to the state of the art:

- We present two novel formalisms that uses Gaussian processes (GP) to predict external temperature and carbon intensity by exploiting (i) periodicities within each single time-series, and (ii) correlations between each time series and another time series (in this case, the temperature at a nearby location and total electricity grid demand) for which a prediction is made available by an outside agency (e.g. a local weather centre and the grid operator). The first extends an existing multi-output GP model that has previously only been used for regression [8], rather than the prediction that we do here. The second uses one time-series to represent the mean value of the second, and applies a single-output GP to model the difference. We empirically evaluate both and indicate their strengths and weaknesses in our domain.

- We present a novel approach to optimising heating whereby the smart controller models the comfort that would result from using a thermostat (set to any particular set-point temperature) to control the heating, and then, optimises heating to ensure the same comfort level, at minimum cost or carbon. We show that this problem has quadratic constraints introduced by the comfort model used, and integer constraints determined by the heater output. Thus, we present both a mixed-integer quadratic program, and a computationally efficient greedy heuristic, to solve it. We show that the greedy heuristic finds solutions that are as competitive as those of a commercial solver (IBM ILOG CPLEX), but in less time.

- We combine these within our smart controller and empirically evaluate it using real UK weather and electricity grid data from January 2010, and show our approach can predict the total cost and carbon emissions over a day to within 9%, and show that over the month it reduces cost and carbon emissions by 15%, and 9%, respectively, compared to a conventional thermostat.

The remainder of this paper is organised as follows. In Section 2 we present our model of the home and heating system. In Section 3 we describe how the smart controller provides real-time feedback, and in Section 4 we describe how it optimises heating use. We conclude in Section 5.

2. HOME THERMAL PROPERTIES

We consider a standard thermal model in which heat leaks from the home (by thermal conduction and ventilation losses) at a rate that is proportional to the temperature difference between the air inside and outside of the home [3]. In more detail, let $\phi \in \mathbb{R}^+$ be the thermal leakage rate of the house measured in W/K. We divide the day into a set of discrete time slots, $t \in T$, and denote the internal temperature of the home at time t as $T_{in}^t \in \mathbb{R}^+$ and the external temperature (in K) as $T_{ext}^t \in \mathbb{R}^+$.

We assume that the home is heated by an electric heat pump, whose thermal output (in kW) is given by r_h. We denote $o_h \subseteq T$ as the set of time slots at which the heat pump is actively producing heat, and also define the variable $\eta_{on}^t \in \{0,1\}$ for every $t \in T$ such that $\eta_{on}^t = 1$ if $t \in o_h$ and 0 otherwise. Given this, the amount of energy delivered (or lost) from the home, η^t, in any time slot is given by:

$$\eta^t = \eta_{on}^t r_h - \phi\left(T_{in}^t - T_{ext}^t\right) + \epsilon^t \tag{1}$$

where ϵ^t is Gaussian noise reflecting additional effects due to the householder's activities (e.g. opening or closing windows, cooking, etc.). The internal temperature of the home after this energy transfer is then given by:

$$T_{in}^{t+1} = T_{in}^t + \frac{\eta^t \Delta t}{c_{air} m_{air}} \tag{2}$$

where Δt is the duration of the time slot (in seconds), and the heat capacity and the total mass of air in the home are $c_{air} \in \mathbb{R}^+$ (in J/kg/K) and $m_{air} \in \mathbb{R}^+$ (in kg) respectively.

The heat pump is controlled by a timer and a thermostat. We denote $o_t \subseteq T$ as the set of time slots at which the heat pump is enabled (but not necessarily actively producing heat), and we define the variable $\eta_{timer}^t \in \{0,1\}$ for every $t \in T$ such that $\eta_{timer}^t = 1$ if $t \in o_t$ and 0 otherwise. The thermostat acts to keep the internal temperature of the home at the thermostat set point, T_{set}, by applying the rule:

$$\eta_{on}^t = \begin{cases} 0 & T_{int}^{t-1} > T_{set} + \Delta T, \\ 1 & T_{int}^{t-1} < T_{set} - \Delta T, \\ \eta_{on}^{t-1} & \text{otherwise.} \end{cases} \tag{3}$$

for all $t \in o_t$. Note that ΔT induces hysteresis such that the thermostat does not continually cycle at the set point.

Finally, we consider the cost and carbon emissions of the electricity that supplies the heater. We assume that the electrical power of the heat pump is given by r_e (in kW). Note that the heat energy produced by a heat pump is greater than the electrical power used, such that, $r_h = COP \times r_e$, where COP, the *coefficient of performance*, is typically between 2 and 4. Thus the total cost of electricity consumed over T is given by $\sum_{t\in o_h} r_e p^t \Delta t$, and the total carbon emissions are given by $\sum_{t\in o_h} r_e c^t \Delta t$, where p^t is the price (pence per kWh), and c^t is the carbon intensity (in gCO_2/kWh), of the electricity at time t.

3. REAL-TIME ENERGY FEEDBACK

Now, to provide real-time feedback to householders on the cost and carbon emissions of their heating preferences, it is not sufficient to simply provide instantaneous measurements of the heater's use, since this requires that the householder extrapolate these figures to the end of the current day. Rather, we require that our smart controller perform this integration automatically, and to do so, it must understand the thermal properties of the home (specifically, the thermal output of the heater and the leakage rate of the home), and also the environmental factors that effect the cost and emissions (specifically, the external temperature and the carbon intensity of the electricity grid). In the next sections, we show how we can learn the former through a simple regression process, and show how we can make prediction of the later, over a 24 hour period, using Gaussian processes.

3.1 Learning Thermal Properties

As described in Section 2, the thermal characteristics of the home is dependent on the thermal output of the heater, the leakage rate of the home, and a number of parameters such as the thermal capacity and the mass of air inside the home. Rather than attempting to estimate all these parameters, the smart controller uses a thermal model, is just as expressive, but which is denoted in terms of the air temperature increase over a time slot, $\mathcal{R} \in \mathbb{R}^+$ (in K), due to the heater and the decrease, $\Phi \in \mathbb{R}^+$ (unitless), due to leakage:

$$\overline{T}_{in}^{t+1} = \overline{T}_{in}^{t} + \mathcal{R}\eta_{on}^{t} - \Phi\left(\overline{T}_{in}^{t} - T_{ext}^{t}\right) \quad (4)$$

Given historical observations (typically over the preceding 24 hours) of internal, T_{in}^t, external, T_{ext}^t, temperatures and the times when the heater was providing heat, η_{on}^t, the smart controller can predict the evolution of the internal temperature over the same period (initialising at $\overline{T}_{in}^1 = T_{in}^1$). The error in this prediction is given by:

$$\sum_{t\in T}\left(\overline{T}_{in}^{t} - T_{in}^{t}\right)^2 \quad (5)$$

Thus, the best estimates of the heater output, $\mathcal{R}$, and leakage rate, Φ, are then those that minimise this error.

3.2 Predicting Environmental Factors

We now turn to predicting the future values of the environmental factors that affect the cost and carbon emissions of the home (specifically, the external temperature and the carbon intensity of the electricity grid). To do so, we exploit the fact that these two time series are likely to be correlated with other time series for which we are able to retrieve predictions over the internet. For example, we expect the external temperature to be correlated to that of the local weather forecast (but not necessarily identical to it due to additional local factors in the immediate environment of the home). Similarly, we observe that carbon intensity is often closely correlated with the total demand for electricity within the grid, for which the grid operators provide accurate predictions in real-time for the next 24 hours. Thus, to predict these correlated time series, we present two novel variations of the Gaussian process (GP): a multi-output GP which explicitly parameterises the correlations between the time series, and a *difference* GP where we model the difference between the normalised time series as a single-output GP.

3.2.1 Gaussian Process Prediction:

A Gaussian process (GP) represents a function as a multivariate Gaussian distribution [9]. It is specified by a covariance function $k(t, t')$, describing correlations between observations of the function at different times, and a mean function, $\mu(t)$, typically taken to be zero, that is the expected value of function prior to seeing any training data. Within the energy domain, Leith et. al used GP for electrical demand forecasting, and shown that it consistently generates significantly more accurate forecasts than comparative time series models [5]. However their work addresses a single time series only, unlike our case where we consider multiple correlated time series.

Now, as discussed above, we have one main time series, for which we have local historical observations, T_M, (i.e. external temperature and carbon intensity), and another correlated time series, for which we again have historical observations, T_C, and also a future prediction over the next 24 hours supplied by an outside agency, T_C^*, (i.e. historical temperature measurements and a 24 hour forecast at the local weather centre, and historical observations of demand and a 24 hour prediction from the grid operator). Our aim is to predict the values of the main time series over the same interval as T_C^*, and we denote these prediction T_M^*.

3.2.2 Multi-Output Gaussian Process:

We first apply a multi-output GP in which the cross-correlation between the time series is explicitly represented as a hyperparameter [8]. To this end, we represent the correlation between two observations in either time series as the Hadamard product of a covariance function over time alone and a covariance function over the time series labels:

$$k([l, t], [l', t']) \triangleq k_L(l, l')k_T(t, t') \quad (6)$$

The covariance function over labels is simply given by:

$$k_L(l, l') \triangleq \begin{cases} 1 & \text{if } l = l', \\ \rho & \text{otherwise,} \end{cases}$$

where $k_L(l, l')$ is unity if both inputs are from the same time series and where ρ is a hyperparameter that determined the cross-correlation between the time series.

The temporal covariance function, k_T, provides more flexibility and is composed of a number of more basic covariance functions. Our first choice is to use the standard *squared exponential* function, given by

$$k_{SE}(t, t') = \sigma_f^2 \exp\left(-\frac{(t - t')^2}{2\ell^2}\right) \quad (7)$$

where σ_f^2 is the amplitude of the process and ℓ is the characteristic length-scale that determines how rapidly the correlation between outputs should decay as t and t' diverge. As the squared exponential implies a strong smoothness assumption regarding the process, we also use the Matérn class

Table 1: Gaussian Process Covariance Functions

Quantity	Method	w	Covariance Function, k_T
Carbon Intensity	Single-output	1	$k_{SE} + k_P + k_N$
	Multiple-output	2	$k_L \cdot (k_M + k_{SE} \cdot k_P) + k_N$
	Difference	2	$k_{SE} + k_{SE} \cdot k_P + k_N$
External Temperature	Single-output	2	$k_{SE} + k_{SE} \cdot k_P + k_N$
	Multiple-output	2	$k_L \cdot (k_M + k_{SE} \cdot k_P) + k_N$
	Difference	1	$k_{SE} + k_{SE} \cdot k_P + k_N$

Table 2: Prediction Accuracy

Quantity	Method	RMSE
Carbon Intensity (gCO_2/kWh)	Single-output GP	32.9
	Multi-output GP	22.3
	Difference GP	20.1
	Linear Regression	27.1
	1-Day Repeat	37.7
	7-Day Repeat	78.4
External Temperature (K)	Single-output GP	2.41
	Multi-output GP	2.13
	Difference GP	2.73
	Weather Forecast	2.36
	1-Day Repeat	3.04

to represent processes with more erratic changes:

$$k_M(t,t') = \sigma_f^2 \left(1 + \frac{\sqrt{3}(t-t')}{\ell}\right) \exp\left(-\frac{\sqrt{3}(t-t')}{\ell}\right) \quad (8)$$

Furthermore, since both the time series considered here exhibit strong daily periodic patterns, we also use a squared exponential periodic covariance function with unit periodicity given by:

$$k_P(t,t') = \sigma_f^2 exp\left(-\frac{sin^2\pi(t-t')}{2\ell^2}\right) \quad (9)$$

Finally, we also consider the noise in measurements by applying an additive gaussian noise function given by:

$$k_N(t,t') = \sigma_n^2 \delta_{t,t'} \quad (10)$$

where δ is the Kronecker delta and σ^2 represents the noise variance.

3.2.3 *Difference Gaussian Process:*

In addition to the multi-output formalism presented above, we also propose an alternative approach, whereby we use a conventional single-output GP to predict the point-wise difference between the normalised outputs of T_M and T_C, over the range of T_C^* to recover T_M^*, and then use the values of T_C^* to recover predictions of T_M^*.

3.2.4 *Empirical Evaluation:*

We evaluate our GP formalisms by comparing them to a number of benchmarks. For carbon intensity, we hypothesise that the value of carbon intensity during any time slot in the future will be equal to the value in that time slot, exactly 1 day, and also 7 days, earlier (reflecting the 1 day and 7 day periodicity observed in the data). We also consider a conventional single-output GP that does not take account of the demand time series, and an approach that determines the correlation between carbon intensity and demand by performing a linear regression between the two, and uses this linear relationship to directly convert a demand prediction into a carbon intensity prediction. In the case of the external temperature, we again consider a single-output GP, the case any time slot is identical to the time slot 1 day earlier, and when we use the online forecast directly.

To perform the evaluation, we use real external temperature sensor data collected from the University of Southampton campus, and also observations and 24 forecasts of Southampton temperature available online through The Weather Channel (`http://uk.weather.com/`). We also make use of carbon intensity and demand observations from the UK grid (for which, again, 24 hour predictions are available - see `http://www.bmreports.com`). All data is at 30 minute intervals, and we perform 31 sequential 24 hour predictions over the entire month of January 2010.

We used Bayesian Monte Carlo to marginalise over the hyperparameters of the GP, and we note that the specification of sensible initialisation values for this has not been actively investigated. We employed a principled approach whereby we investigated the autocorrelation properties of the time series under consideration. The knowledge of such correlations also helped guide the choice of appropriate covariance structure. We set the history length for best performance based on empirical observations, and Table 1 shows the covariance functions, and the number of days of historical data, w, used to generate the predictions in this section.

Our results, in terms of the root mean square error of these predictions, are shown in Table 2. We note that the multi-output GP significantly outperform our benchmark approaches on both prediction tasks. This approach is able to effectively exploit both the periodicities within each single time series, and also correlations between them. We also note the while the difference GP works very well in the case of carbon intensity (providing the best predictions overall), it performs less well on external temperature. When the time series are strongly correlated, both approaches provide good predictions. However, when correlations are not present (as is sometimes the case in the external temperature time series) the difference GP has no mechanism to compensate. In contrast, the multi-output GP can adjust the hyperparameter that describes this correlation and perform well. Thus, the difference GP can be an effective approach, but must be used with care since it is not adaptive to changing correlation.

3.3 Estimating Cost and Carbon Emissions

Now, given the external temperature and carbon intensity time series predicted above, the thermal properties of the home learned earlier, and the description of the thermostat, the smart controller can predict the operation of the heater over the course of the day, and hence predict the total cost and carbon emissions that will result from any particular thermostat setting. To evaluate these predictions we have simulated a home and the smart controller installed within it (see Figure 2). In doing so, we use leakage rate $\phi = 90$ W/K and $m_{air} = 1205$ Kg (corresponding to a small well-

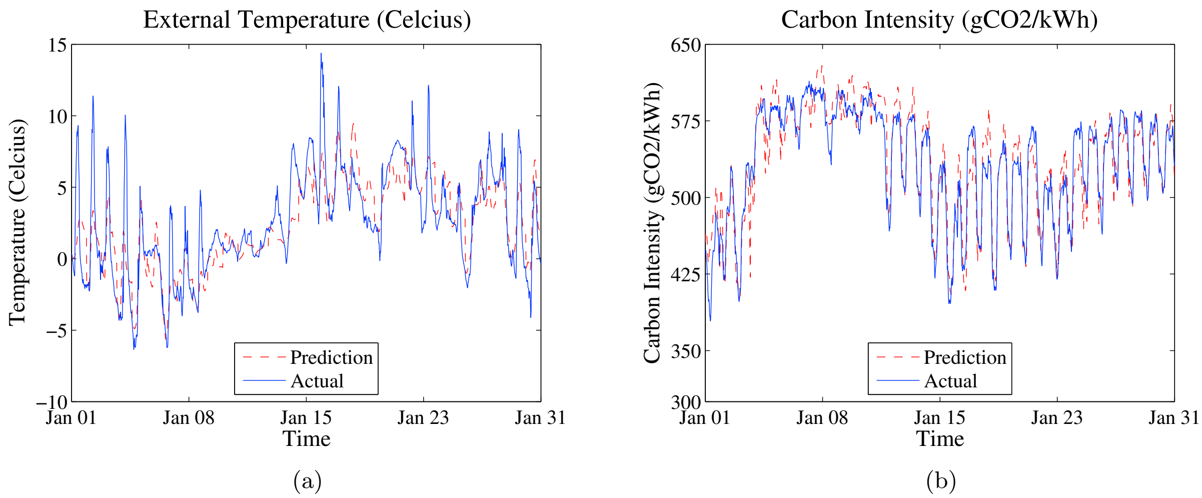

Figure 1: Comparison of Gaussian process predictions for (a) external temperature and (b) carbon intensity using real data for January 2010.

insulated home), an electric heat pump with a thermal heat output of 2.5kW and COP of 2.5, and used a standard value of $c_{air} = 1000$ J/Kg/K in the thermal model. We consider a thermostat set point, $T_{set} = 20$ C, and we assume that the heater is enabled (but not necessarily actively producing heat) between 07:00 and 23:00 each day. Furthermore, we consider a setting in which the price of electricity exhibits a critical pricing period between 15:00 to 20:00 when the price of electricity is £0.24 per kWh, a cheap overnight rate between 20:00 and 06:00 of £0.08 per kWh, and a standard rate of £0.12 per kWh at all other times.

Using the same data for January 2010, we predict the cost and carbon emission for each day, comparing these to the actual values as the true external temperature and carbon intensity is revealed. Over this period, the actual mean daily heating cost is £1.59 and the root mean squared error in our predictions is £0.14, while the actual mean daily carbon emissions are 5.99 kg and the root mean squared error in our predictions is 0.54 kg (an error of 9% in both cases). We observe that in both cases, the prediction errors are less those achieved by simply using the current day as a prediction for the next (in this case, £0.16 and 0.63 kg, respectively). Crucially, our approach allows the effects of changes to the thermostat set-point, or heater timing, to be calculated since it explicitly models these variables. Furthermore, we note that the maximum error occurs at the very start of the day; updating the predictions over the course of the day will cause the predicted cost and carbon to converge to the actual cost and carbon by the end of the day. Figure 2 shows this cost information being displayed on the smart controller's graphical user interface.

4. HOME ENERGY OPTMISATION

Having shown how our smart controller can predict the cost and carbon emissions of the home's heating system, we now show how it can adapt to real-time cost and carbon intensity signals, adjusting the timing of heater use in order to satisfy the home owner's preferences for comfort minimising cost and carbon emissions. Our approach uses mathematical programming to ensure that the smart controller delivers the same level of comfort as a standard thermostat whilst also minimising either cost or carbon. This is attractive since it does not require the householder to explicitly trade-off between cost and comfort (or carbon and comfort), nor indeed, even be aware of the underlying comfort model.

4.1 Evaluating Comfort

We use a standard model of comfort, based on the ANSI / ASHRAE Standard 55-2010, that has previously been used within the smart home environment [3]. To this end, we denote $o_c \subseteq T$ as the set of time slots at which the householder requires comfort, and we define the *instantaneous* discomfort, $\Delta d^t \in \mathbb{R}^+$, such that:

$$\Delta d^t = \begin{cases} \omega_1(\tau^t_{in} - T_{opt})^{\phi_1} & T^t_{in} \geq T_{opt}, \\ \omega_2(T^t_{in} - T_{opt})^{\phi_2} & T^t_{in} < T_{opt} \end{cases} \tag{11}$$

Then, the actual discomfort at time t is a combination of the instantaneous discomforts at t and at $t-1$, such that:

$$d^t = \Delta d^t + \gamma \Delta d^{t-1} \tag{12}$$

where $\gamma \in [0, 1]$ scales the effect of the previous time slot on the current one (capturing the psychological persistence of discomfort). The total discomfort over is the sum of the discomfort at every $t \in o_c$ given by given by $\sum_{t \in o_c} d^t$.

4.2 Optimising Heating

Given this comfort model, the smart controller can predict the total discomfort that the thermostat will to deliver over the day. We denote this as the target discomfort, D_{target}. The aim of the smart controller is then to determine when the heater should on in order that the total discomfort is less than the target discomfort with the minimum cost or carbon emissions. More formally, specifying $\eta^t_{on} \in \{0, 1\}, \forall t \in T$ as the decision variables, the objective function for cost and carbon respectively, is given by:

$$\arg\min_{\eta^t_{on}} \sum_{t \in T} \eta^t_{on} r_e p^t \Delta t \qquad \text{and} \qquad \arg\min_{\eta^t_{on}} \sum_{t \in T} \eta^t_{on} r_e c^t \Delta t$$

where the evolution of the internal temperature is described

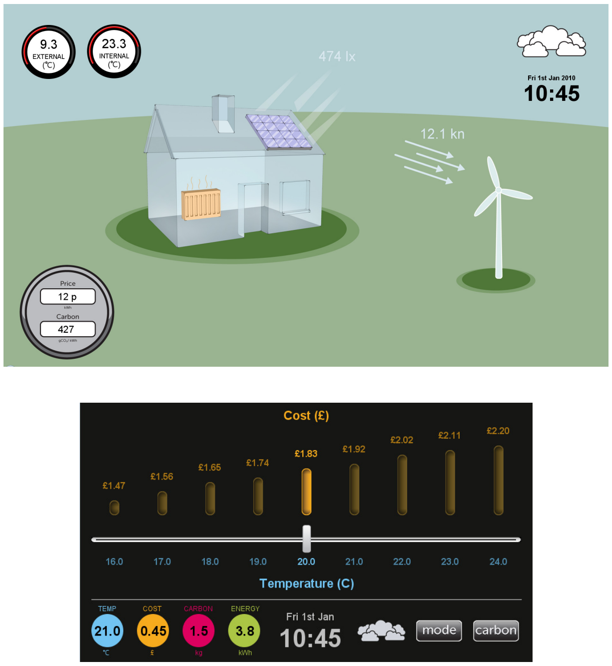

Figure 2: Simulation of the home, the smart controller and its graphical user interface.

by Equation 4, that of discomfort by Equations 11 and 12, and the constraint that $\sum_{t \in o_c} d^t \leq D_{target}$.

4.2.1 *Mixed-Integer Quadratic Programming:*

Now, it is typical to consider that $\phi_1 = \phi_2 = 2$ within Equation 11, and thus the objective function is quadratic. In this case, the optimisation described above can be directly implemented as a mixed-integer quadratic program, using a standard solver, such as IBM ILOG CPLEX.

4.2.2 *Greedy Heuristic Optimisation:*

More generally, if $\phi_1 \neq \phi_2 \neq 2$, or if the computational resources to implement a full mixed-integer quadratic solver are not available on the hardware installed within the home, we also present a greedy heuristic (see Algorithm 1) that incrementally finds the individual time slot where switching the heater on results in the largest reduction in discomfort, for the minimum additional cost or carbon. This process repeats until the total discomfort in equal to, or below, that of the target discomfort.

Algorithm 1 Greedy Algorithm to Optimise Home Heating

// Calculate the initial discomfort level

for $j = 1$ to $|T|$ **do**
  $\eta^j_{op} \leftarrow 0$
  Compute T^t_i, Δd^t_{test} and d^t_{test} as per (4), (11) and (12).
end for
$D_{op} = \sum_{t \in o_c} d^t_{test}$

// Loop through until discomfort is below target

while $D_{op} > D_{target}$ **do**
  $index \leftarrow 0$
  $bestratio \leftarrow 0$
  for $j = 1$ to $|T|$ **do**
    $\eta_{\mathbf{test}} = \eta_{\mathbf{op}}$
    $\eta^j_{test} = 1$
    for $t = 1$ to $|T|$ **do**
      Compute T^t_i, Δd^t_{test} and d^t_{test} as per (4), (11) and (12).
    end for
    $D^j_{test} = \sum_{t \in o_c} d^t_{test}$
    $\Delta D = D_{op} - D^j_{test}$
    if $\frac{\Delta D}{p_t} \geq bestratio$ **then**
      $bestratio = \frac{\Delta D}{p_t}$
      $index \leftarrow j$
    end if
  end for
  if $index > 0$ **then**
    $\eta^{index}_{op} \leftarrow 1$
    $D_{op} \leftarrow D^j_{test}$
  end if
end while

4.3 Empirical Evaluation

We now compare the performance of the mixed-integer quadratic solver and the greedy heuristic against thermostatic control within the simulation setting previously described. For all experiments, comfort is required between 8am and 11pm, and we use values of $\omega_1 = \omega_2 = 1$, $\phi_1 = \phi_2 = 2$, $\gamma = 0.8$ and $T_{opt} = 23.5$ °C in the comfort model. We consider time slots that are 5 minute long, such that $|T| = 288$, we perform 31 sequential optimisations over the month of January 2010. We consider 5 minutes time slots (i.e. $|T| = 288$), and thus, we limit the mixed-integer quadratic solver to 5 minutes of running time on a standard desktop PC (the greedy heuristic runs in milliseconds).

Figure 3 show examples of the optimisation process. In this case, using the MIQP optimisation routine to minimise both cost and carbon. In both cases, we present the internal temperature of the home when the heating is controlled by both the standard thermostat and the smart controller. The green shaded area represents the time interval over which comfort is required, and the red shaded area represents when the heating system is actually producing heat.

Consider Figure 3(a) where the smart controller is minimising cost. Note that the smart controller applies heat before the peak pricing period, allowing the temperature to increase, and then allows this heat to leak away over this period. In contrast, the standard thermostat applies heat uniformly across this period. Similarly, note that the smart controller also exploits the low price of electricity before 06:00 and supplies heat even though it is not immediately required. In both cases, it is effectively storing cheap electricity in the form of hot air, so that this stored energy can be used when electricity is more expensive, and this approach provides an alternative to the use of more costly electrical storage batteries proposed by [11].

Figure 3(b) where the smart controller is minimising carbon. As before, we compare the internal temperature of the home when the heating is controlled by both the standard thermostat and the smart controller. Again the smart controller applies more heat before the peak period when carbon emissions would be the highest, allowing the temperature to increase, and then allows this heat to leak away over this period. In contrast, the standard thermostat again applies heat uniformly across this period (indeed, the operation of

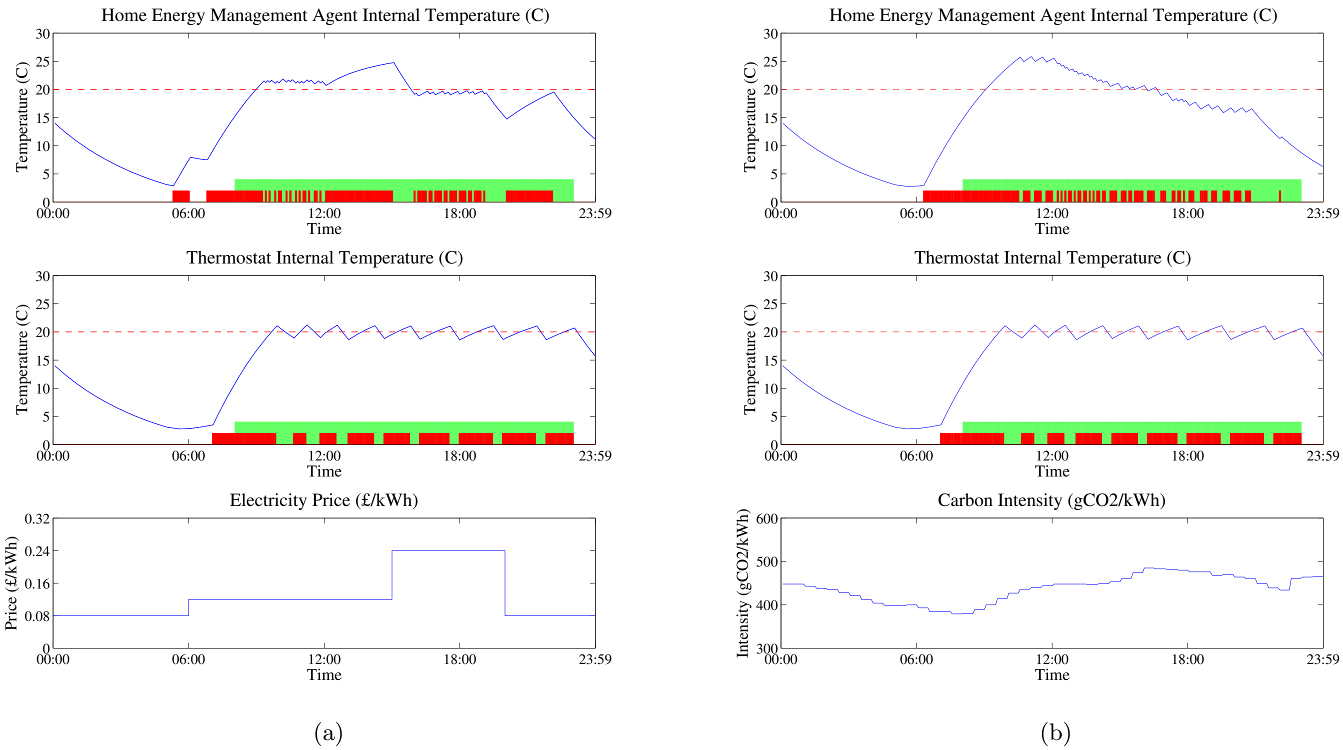

Figure 3: Example showing the MIQP formulation optimising heating use to maintain comfort whilst (a) minimising cost in a setting with a critical price period between 15:00 and 20:00, and (b) minimising carbon emissions given the variable carbon intensity of the electricity grid.

.

the thermostat is identical to the previous case). In addition, note that the smart controller also exploits the lower carbon estimates before 06:00 and supplies heat even though it is not immediately required.

Finally, in Table 3 we compare the total cost and carbon emissions of both the MIQP and greedy heuristic over the month of January 2010. Overall, the smart controller reduces cost and carbon emissions by 15%, and 9%, respectively, compared to a standard thermostat, and that the greedy heuristic performs well, being within 1% of the mixed-integer quadratic program (MIQP) solutions (at a fraction of the computational cost).

5. CONCLUSIONS

In this paper, we addressed the challenge of adaptively controlling a home heating system in order to minimise cost and carbon emissions in response to real-time price and carbon intensity signals. We proposed a home energy management agent, or smart controller, that operated on behalf of the householder, and learned the thermal properties of the home, used a Gaussian process model to predict the local environmental parameters, in order to both provide real-time feedback concerning the cost and carbon emissions of their heating preferences, and optimised their energy use. We showed, using real-world data and a simulated home, that our approach could reduce cost and carbon emissions by up to 15%, and 9%, respectively compared to using a standard thermostat.

Our future work is focused on deploying the home energy management agent described here within a number of real homes in collaboration with our industrial partners. The key challenges that we expect to face in doing so, are ensuring that the relatively compuationally expensive prediction and optimisation processes can be performed effectively on the low power devices to be deployed, and to ensure that the thermal model that we have developed is sufficiently representative of a real house, such that it is able to accurate costs and carbon emissions with the necessary accuracy.

Table 3: Comparison of smart and thermostat control when optimising for cost and carbon emissions for January 2010.

	MIQP	Greedy	Thermostat
Cost (£)	42.44	42.81	50.37
Carbon (kg CO_2)	174.60	174.90	191.48

6. ACKNOWLEDGEMENTS

The work presented in this paper was joint funded the iDEaS project (`www.ideasproject.info`), and through the Intelligent Agents for Home Energy Management project (EPSRC EP/I000143/1).

7. REFERENCES

[1] DECC. Smarter grids: The opportunity. Technical report, Department of Energy and Climate Change

(DECC), 2009.

[2] DOE. *Grid 2030: A National Vision For Electricity's Second 100 Years*. United States of America Department of Energy, 2003.

[3] Y. Guo, R. Li, G. Poulton, and A. Zeman. A Simulator for Self-Adaptive Energy Demand Management. In *Proceedings of the Second IEEE International Conference on Self-Adaptive and Self-Organizing Systems*, pages 64–73, 2008.

[4] Kempton. Two theories of home heat control. *Cognitive Science*, 10(1):75–90, 1986.

[5] D. Leith, M. Heidl, and J. Ringwood. Gaussian process prior models for electrical load forecasting. In *Proceedings of the International Conference on Probabilistic Methods Applied to Power Systems*, pages 112 –117, 2004.

[6] Low Carbon. The UK low carbon transition plan: National strategy for climate and energy. Technical report, Department of Energy and Climate Change (DECC), 2009.

[7] A. Molderink, V. Bakker, M. Bosman, J. Hurink, and G. Smit. Management and control of domestic smart grid technology. *IEEE Transactions on Smart Grid*, 1(2):109–119, 2010.

[8] M. A. Osborne, S. J. Roberts, A. Rogers, S. D. Ramchurn, and N. R. Jennings. Towards real-time information processing of sensor network data using computationally efficient multi-output gaussian processes. In *Proceedings of the Seventh International Conference on Information Processing in Sensor Networks*, pages 109–120, Washington, DC, USA, 2008.

[9] C. E. Rasmussen and C. K. I. Williams. *Gaussian Processes for Machine Learning (Adaptive Computation and Machine Learning)*. The MIT Press, 2005.

[10] F. Schweppe, B. Daryanian, and R. Tabors. Algorithms for a spot price responding residential load controller. *IEEE Power Engineering Review*, 9(5):49 – 50, 1989.

[11] P. Vytelingum, T. D. Voice, S. D. Ramchurn, A. Rogers, and N. R. Jennings. Agent-based micro-strorage management for the smart grid. In *Proceedings of the Nineth International Conference on Autonomous Agents And Multiagent Systems*, pages 39–46, Toronto, Canada, 2010.