Lies, Damn Lies and Preferences; a Gaussian Process Model for Ubiquitous Thermal Preference Trials.

Damien Fay, Kenneth N. Brown, Liam O'Toole
Cork Constraint Computation Centre, Department of Computer Science
University College Cork, Ireland
d.fay@4c.ucc.ie

Abstract

This paper presents a study of user comfort levels using an ubiquitous interface. The aim is to analyse the comfort function of an *individual* as opposed to previous approaches that look at the average human being. The data is analysed using Gaussian Process regression which allows several mechanisms to be exploited. These include regression on the data to give an estimate of a users comfort function. The prediction variance is also estimated and outlier influence can be reduced easily. In addition, a natural means of combining the preferences of users falls out of the approach. The combination algorithm takes into account fairness tempered by the quality of the user' preference estimates. Empirical results show that the combined preferences have a well defined maxima which can be used as a control signal for a HVAC system. The Gaussian Process approach is hierarchical and interestingly, while those users studied have differing preferences, their hyperparameters (at the second level of the hierarchy) are concentrated; i.e. there is a strong commonality across individuals in this domain.

General Terms

Categories and Subject Descriptors

G.3 [**Probability and statistics**]: Correlation and regression analysis; H.1.2 [**Information Systems**]: User/Machine Systems

Keywords

User preference, Gaussian process, ASHRAE, PMV.

1 Introduction

Energy usage in buildings accounts for between 20% and 40% [15] of total energy consumption, with Heating Ventilation and Air Conditioning (HVAC) systems accounting for 50% of this figure in the USA. While much effort has been focused on developing efficient HVAC systems, including efficient control, less attention has gone into investigating the setpoints maintained by these systems. Typically the setpoints are fixed or adjusted manually for different seasons, regardless of the preferences of the occupants of the building. In this paper we present a Bayesian technique for estimating a user's thermal comfort as a function of their *current* environment, based on a history of their reported perceptions of comfort.

Thermal comfort sampling and modelling is a particularly difficult task. There are many factors which effect a user's perception of comfort, including outside temperature, the internal environment, the clothes they are wearing, their type of activity (metabolic rate), the mood they are in, their current health and many more. The largest set of studies on user comfort levels comes from the ASHRAE RP-884 Database [5], which contains 22 thousand responses from 20 different studies undertaken mainly during the 1990s in which differing levels of detail were recorded, including clothing index, metabolic rate and the users perception, via questionnaires. We build on this research but use a different paradigm made possible by ubiquitous computing. Our data is gathered from three experiments with 65 volunteers. We take real-time comfort measurements via a computer interface in a work environment where the full attention of the user is not sought. Our aim is to produce a comfort indicator for the *individual* as opposed to one for the average occupant. While sensors can be placed near a user these cannot measure every factor, so our aim is to include only variables which can reasonably be measured (we choose temperature and humidity) and to produce an effective estimate from them. The other unmeasured factors thus feed into the data as noise. However, there is another source of error, on the dependent variable side of the function: the users themselves are not a good measurement device for their own comfort. A user may not be aware that it is warmer than they would normally prefer. In addition, due to annoyance or engagement with a more important task, users can respond with outliers. Thus our software interface is designed to minimise the nuisance to the user. Our study includes an analysis of the influence of the data gathering tool on the users' responses.

The specific statistical approach used in the study is based on Gaussian Process (GP) Regression. While involved, the GP approach solves the aforementioned problems allowing a good estimate of a user's preference function and a natural way in which to combine the preference functions of individuals sharing the same office space. We apply our GP

Buildsys'12, November 6, 2012, Toronto, ON, Canada.

approach to our experimental data, and we show how the the GP can incorporate belief about the measurement noise. We demonstrate a variance reduction approach to reduce the effect of outliers. To the best of our knowledge Gaussian Processes have not been applied in this area before.

2 Related Work

User thermal preferences have been traditionally defined by the Ashrae Standard [5]. The standard defines thermal comfort using the *Predicted Mean Vote* (PMV) [7], a complex recursive equation which estimates a user's response on a 7 point scale (see Figure 1) using a full array of environmental factors including clothing index and metabolic rate. In practice however, many of these factors are not measured and so their application can be broken into two types: fixed setpoint and adaptive setpoints [5]. In the fixed setpoint approach a fixed setpoint/preference is given based on the average environment for a building/space of this type. In Ireland, for example, air conditioned offices have a fixed setpoint of 21°C in Summer and 24°C in Winter. However, given live readings of some environmental conditions, the adaptive approach can be employed. In this approach variables such as the clothing index and metabolic rate are fixed to their average while those that can be measured are inserted into the equations and the temperature setpoint that optimises Predicted Mean Vote then adapts to the environment continually [20] [14] [2]. The PMV itself can take a significant amount of resources to estimate and so is sometimes modelled using a neural network [2] or fuzzy logic system [6]. However, the adaptive approach does not consider the *individual's* thermal preference. Perhaps the closest research in spirit to the current research is that by Daum et al. [4]. In their study, a pop-up polled users four times a day about their perceived comfort levels, their clothing index and their activity levels. Our approach polls only for perceived comfort level as we consider the other information too large an intrusion on users activities and the emphasis here is on the quality of the data returned in a ubiquitous approach. In addition, the modelling procedure employed by Daum et al. is logistic regression in which the data is quantized. Gaussian processes have become increasingly popular for data analysis due to their flexibility and the excellent results they produce [16]. They have been applied in many fields such as robot control [12] and fault detection [19], as a general optimization approach [1][17][9][18].

3 Interface and data collection procedure.

The user interface in this study is designed to be easy to use and cause minimal intrusion. Figure 1 shows a cropped screenshot in which the user interface appears in the top left corner of the user's computer screen. There is a colorbar, with 7 labels to the right (hot, warm, neutral etc.). Once a click has been received the interface disappears and does not appear again until a random time the next day. The deviation from neutral is the dependent variable in the model and is denoted Y, and lies in the range -133 to 133 (the height of the colorbar in pixels is 133×2).[1] Although the labels in

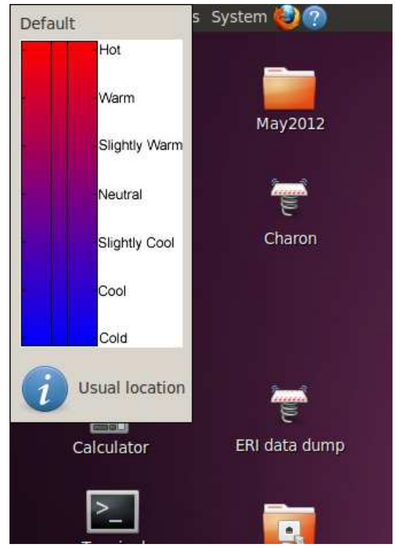

Figure 1. A screenshot of the user interface (top left).

Figure 1 correspond to the ASHRAE labels there is an important difference: the ASHRAE polls consists of 7 discrete levels, while here they are continuous, which allows us to avoid quantisation errors.

The preference collection software was deployed in three different sites. Table 1 summarises the data collected. The first is at a room in University College Cork (UCC) in which each user has a temperature and humidity sensor located at their desk. This site is the most important in terms of instrumentation and control but contains the fewest participants: just four. The HVAC system consists of underfloor storage heating and natural ventilation; the room also has a large south facing window. The environment is difficult to control, as there is a lag of approximately eight hours between the control/storage heating and the resulting temperature in the room, and this site presents the largest variation in the environment of the three.

The second implementation took place at the new engineering building at the National University of Ireland, Galway (NUIG). This building is in its first year of operation. The building contains 2 large rectangular open plan research rooms, each containing space for approximately one hundred students. Both rooms have large south facing windows. From these two rooms forty participants were recruited. Each room consists of 5 temperature and CO_2 sensors equally spaced along the room's length.

The final implementation was in a research building at Cork Institute of Technology (CIT). These rooms have temperature sensors in various locations: for the 21 participants there are 13 sensors. The HVAC system is modern but has been in use for several years and so the environment showed the least variation.

In addition hourly readings of the external temperature are available for each site.

[1] The direction of the data might cause confusion; in keeping with image processing convention the image begins at the top-left corner and so while Neutral is zero, Hot is -133 and not +133. While it would have been easy to reverse this we found that doing so makes the 3-D figures in this paper harder to interpret.

Table 1. Summary of data collected.

Site	No. Participants	Average number of polls (Max,Min,Std)	Internal Temperature	Internal Humidity	External Temperature	Environment
UCC	4	39 (60,4,15)	Yes	Yes	Yes	Controlled but large variation.
NUIG	40	39 (60,4,15)	Yes	No	Yes	Controlled but not fully tuned.
CIT	21	39 (60,4,15)	Yes	No	Yes	Controlled and tuned.

4 Gaussian processes.

In the GP's presented here, the variables of interest are updated using Bayesian inference. A *prior* belief is first given about the process of interest. Then, given new information (i.e. a poll), this belief is updated, resulting in a *posterior*. By *belief* we mean a probability distribution and *all* unknowns in a process are assigned a (probability) distribution. Thus, unlike classical statistical modelling, Bayesian modelling does not force us to accept one particular value for a parameter but rather accepts all values as true with differing levels of probability. In the current setting, this belief is the comfort of the individual at the current environmental settings, based on their past responses.

4.1 Gaussian Process Models

The data model used is a GP [17][16]. A GP consists of assigning a kernel to each measured data point and performing a regression based on that kernel and the measurement noise at that data point. While most regression techniques use a single function to fit the data, GPs are equivalent to taking a weighted average of functions over a whole functional space; thus reducing the dependence on a single function and in addition providing an estimate of the variance of the mean estimates.

The first stage in the GP employed here is input data scaling. Given a d dimensional input, $x \in \Re^d$, the data is first scaled such that the variation in all dimensions along the data is the same. To achieve this the data is typically multiplied by a scaling matrix $M \in \Re^{d\times d}$; which in this case is diagonal. M thus consists of d elements and we choose to keep the first entry on the diagonal as 1. The reason for this choice is that the first variable is internal temperature, the dominant variable, which is already in units that can be easily interpreted. Thus M consists of $d-1$ *unknown* elements denoted, $\{m_1,\ldots,m_{d-1}\}$. These unknowns can be incorporated into the set of hyperparameters for the GP. The estimation procedure is discussed below.

A GP is defined as a process in which realisations from the process are jointly multivariate Normally distributed. Specifically, the data generated by the process at n sample points, $Y_{x_{1:n}}$, are drawn from a multivariate Gaussian distribution as:

$$Y_{x_{1:n}} \sim \mathcal{N}[\mu_x, C_{\mathbf{x},\mathbf{x}}] \tag{1}$$

where $\mathcal{N}$ denotes a Gaussian distribution, $x_{1:n}$, denotes n samples taken at points $x_1 \ldots x_n$, $\mu_x \in \Re^d$, is the mean of the process and $C_{\mathbf{x},\mathbf{x}}$ is the covariance matrix. The covariance matrix depends on the separation of the data points and so is often referred to as the covariance function. Following appropriate scaling of the inputs (via M) an **isotropic** covariance function can be used in which the variation of the function is equal in all directions. Given an isotropic covariance function it now becomes more convenient to talk in terms of the correlation function which is related to the covariance function via[2]:

$$C_{\mathbf{x},\mathbf{x}} = \sigma_x^2 R_{\mathbf{x},\mathbf{x}} \tag{2}$$

as σ_x^2, is the variance of the process and $R_{\mathbf{x},\mathbf{x}} \in \Re^{n\times n}$ is the correlation function. The correlation function is defined by a kernel with the following properties:

$$R_{x,x} = R(0) = 1 \tag{3}$$

the correlation function at a distance of zero (i.e. between a point x and itself) is one,

$$R_{x_1,x_2} = R(\|x_1 - x_2\|) \tag{4}$$

the correlation function is only a function of the separation of the sample points, where $\| \bullet \|$ denotes euclidean distance, and

$$R(h_1) < R(h_2) \text{ for } h_1 > h_2 > 0 \tag{5}$$

the correlation function dies away as the distance increases; where h denotes a distance. This last condition is necessary in order for single function paths (i.e. a collection of samples taken from the GP) to be used for inference [17]. A valid kernel function may be derived from any symmetric probability density function (pdf). In this application, the Matérn Kernel, derived from the t-distribution is used as it has just two parameters and allows the kernel to have a variety of shapes, from a Gaussian-like shape to shapes peaked at zero *and* with a long tail. The Matérn kernel is defined as:

$$R(h,\theta,\nu) = \frac{1}{\Gamma(\nu)2^{\nu-1}}\left(\frac{2\sqrt{\nu}|h|}{\theta}\right)^{\nu} K_\nu\left(\frac{2\sqrt{\nu}|h|}{\theta}\right) \tag{6}$$

where K_ν is the modified Bessel function, θ and ν are parameters of the kernel with θ controlling the scale and ν the shape of the kernel.

Now, given a set of points at which samples have already been taken, $x_{1:n}$, and a set of locations (called evaluation points), x^*, at which we have not sampled, the relationship between the sampled and evaluation points may be expressed by partitioning Equation 16 in terms of the cross and auto-correlation matrices of the sampled and evaluation points as [17]:

$$\begin{bmatrix} Y_{x_{1:n}} \\ Y_{x^*} \end{bmatrix} \sim \mathcal{N}\left[\begin{bmatrix} \mathbf{1_n} \\ \mathbf{1}_* \end{bmatrix}\mu_x, \sigma_x^2 \begin{bmatrix} R_{\mathbf{x},\mathbf{x}} & R_{\mathbf{x},x^*} \\ R_{\mathbf{x},x^*}^T & R_{x^*,x^*} \end{bmatrix}\right] \tag{7}$$

where Y_{x^*} is the value of the process at x^*, $\mathbf{1_n}$ and $\mathbf{1}_*$ are appropriately dimensioned vectors of ones, $R_{\mathbf{x},\mathbf{x}}$ is the auto-correlation between the known sample points, $R_{\mathbf{x},x^*}$ is the cross-correlation between the sample and evaluation points

[2]Here is assumed that the overall process mean is zero; alternatively a non-zero mean may be subtracted from the data prior to modelling.

and R_{x^*,x^*} is the auto-correlation of the evaluation points. $\sim$ denotes *drawn from* and $\mathcal{N}$ denotes a Gaussian distribution.

The unknown parameters in the GP defined by Equation 7 are the process mean, the process variance and the shape parameters for the kernel. In addition, the data scaling parameters are unknown and so the full set of unknowns is $\{\mu_x, \sigma_x^2, \theta, \nu, m_1, \ldots, m_{d-1}\}$. These may be estimated iteratively using Bayesian conjugate analysis and a hierarchical GP in which the parameters are organised in a hierarchy as:

$$[\mu_x, \sigma_x^2, \theta, \nu] = [\mu_x|\sigma_x^2] \times [\sigma_x^2] \times [\theta, \nu, m_1, \ldots, m_{d-1}] \tag{8}$$

where $[\bullet]$ denotes distribution. It is thus assumed that the kernel and data scaling parameters are independent of the process mean and variance. They are estimated first, followed by the next stage in the hierarchy; estimating the variance. The process mean is then estimated (conditional on the variance). Finally, estimates of the function at particular sample points may be made given the hyperparameters. Each stage is now explained further.

Given the current data points the kernel and data scaling parameters are estimated by maximising the log likelihood function with respect to $[\theta, \nu, m_1, \ldots, m_{d-1}]$. This function (the log marginal likelihood) may be expressed as:

$$L = -\frac{1}{2} Y_{x_{1:n}}^T C^{*-1} Y_{x_{1:n}} - \frac{1}{2} log|C^*| - \frac{n}{2} log(2\pi) \tag{9}$$

where $L \equiv \log p(Y_{x_{1:n}}|x_{1:n}, \theta, \nu, m_1, \ldots, m_{d-1})$ is the log likelihood function, $C^* = \sigma_x^2(R_{\mathbf{x},x^*} + \zeta_x)$, is the covariance matrix of the noisy data and ζ_x is the measurement noise covariance. Equation 9 has two terms on the right hand side which are quite interesting. The first term, $Y_{x_{1:n}}^T C^{*-1} Y_{x_{1:n}}$ is the data fit while the second term, $log|C^*|$, is the variance of the model itself (the third term is just a normalisation factor). Thus there is a natural trade-off between the data fit and the model complexity which avoids over fitting (see [16]). The process parameters are estimated using a conjugate Bayesian approach in which the standard conjugate Bayesian prior for a mean and (unknown) variance are used; specifically the mean has a normal prior; $\mu_x \sim \mathcal{N}\left[0, \sigma_x^2 \delta^2\right]$ and the variance has a normal-inverse-Gamma prior; $\sigma_x \sim \mathcal{IG}\left[a/2, b/2\right]$. δ is our initial estimate of the variance of the mean. The variance is initially assumed to lie in the interval $[a, b]$.

An estimate of the function at the evaluation points may be constructed using least squares (see [17] for details). However, in the current application we are interested in estimating the value of Y_{x^*} given that the measurements are noisy. In the presence of measurement noise Equation 7 becomes [16]:

$$\begin{bmatrix} Y_{x_{1:n}} \\ Y_{x^*} \end{bmatrix} \sim \mathcal{N}\left[\begin{bmatrix} \mathbf{1_n} \\ \mathbf{1}_* \end{bmatrix} \mu_x, \sigma_x^2 \begin{bmatrix} R_{\mathbf{x},\mathbf{x}} + \zeta_x & R_{\mathbf{x},x^*} \\ R_{\mathbf{x},x^*}^T & R_{x^*,x^*} \end{bmatrix}\right] \tag{10}$$

and an estimate of the value of the function at the evaluation points, x^* may be expressed as [16]:

$$\hat{Y}_{x^*} = \mathbf{1}_* \hat{\mu}_x + R_{\mathbf{x},x^*} \left(R_{\mathbf{x}} + \zeta_x\right)^{-1} \left(Y_{x_{1:n}} - \mathbf{1_n} \hat{\mu}_x\right) \tag{11}$$

An estimate of variance at the evaluation points, $\sigma^2(x^*)$, may be estimated via [16]:

$$\hat{\sigma}^2(x^*) = \hat{\sigma}_x^2 \left(R_{x^*} - R_{\mathbf{x},x^*}^T K R_{\mathbf{x},x^*} + \frac{\left(1 - \mathbf{1_n}^T K R_{\mathbf{x},x^*}\right)^2}{\mathbf{1_n}^T K \mathbf{1_n} + \delta^{-2}}\right) \tag{12}$$

where $K = (R_{\mathbf{x},\mathbf{x}} + \zeta_x)^{-1}$ is used to simplify notation. The maximum *a-posteriori* estimates for the process parameters are [12]:

$$\hat{\mu}_x = \left(1 - K + \delta^{-2}\right)^{-1} \mathbf{1_n}^T K Y_{x_{1:n}} \tag{13}$$

and

$$\hat{\sigma}_x^2 = \frac{\left(b + Y_{x_{1:n}}^T K Y_{x_{1:n}} - \left(\mathbf{1_n}^T K \mathbf{1_n} + \delta^{-2}\right)\right) \hat{\mu}_x^2}{n + a + 2} \tag{14}$$

4.2 Outlier influence reduction.

At this stage the base model has been presented but there still remains one set of unknowns in the equations above, the measurement noise at the sample points, ζ_x. In many cases this is assumed to be constant [16] or another GP is employed to model the variance in addition to the mean (see [8] [11]). However here we use an approach similar to that in [18]. In the current setting, an initial guess can be made for the variance of a sample/poll (see Section 5.2.1), say ζ_x^0. Given this initial guess the GP can be used to produce an estimate of the process at x, i.e. $\hat{Y}_x$, and this can be taken from the measured value to produce a residual. A recursive procedure can then be used to further estimate the variance based on the residual as:

$$\zeta_x^i = \alpha \zeta_x^{i-1} + (1 - \alpha) r_x^i \tag{15}$$

where $r_x^i = \hat{Y}_x - Y_x$ is the residual at x in iteration i and α is a coefficient which is here set to 0.7 (although this algorithm was found to be robust to different values of α). The net effect of this algorithm is that outliers are excluded as shown empirically in Section 5.2.1.

5 Results

The results are broken into three major sections; the first presents a preliminary analysis of the data demonstrating major characteristics (5.1). The aim of the second section is to show the GP model in operation (5.2). The third section concentrates on the group results from the three sites (5.3).

5.1 Preliminary data analysis

Figure 2 shows the responses of a typical user with respect to the internal temperature (note that there are other dimensions to the data not shown here). As can be seen the relationship between internal temperature and the deviation, Y, is not immeadiatly clear. In addition, there appears to be a band around $Y = 0$ for which the response is invariant to the input. This band is called the *deadband* and shows a clear band of entries from 21-26 °C for which the user gave a near zero response. The likely cause of this deadband is several fold. It is well known that biases exist in survey responses known as the *demand characteristics* [13]. Also, anecdotally, users indicated that when they didn't care they usually clicked at the neutral position. In addition, skin is very bad at perceiving warmth and cold [10] especially when the room

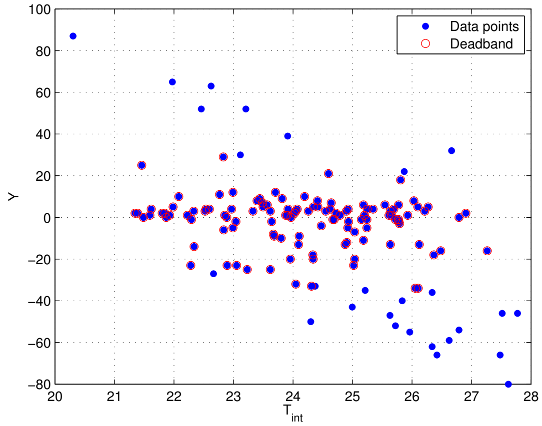

Figure 2. The response of a single user with respect to internal temperature; the deadband location.

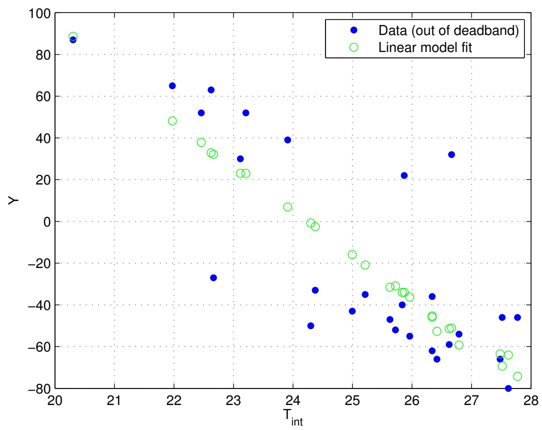

Figure 3. The response of a single user with respect to internal temperature; Points outside the deadband and a linear fit to those points.

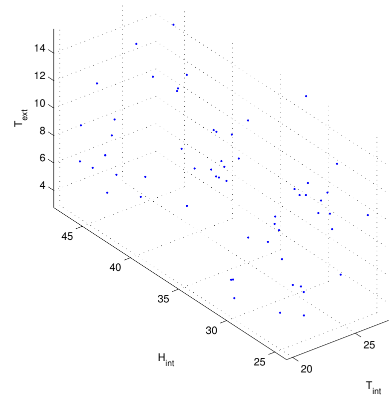

Figure 4. Original input data points; note the axis scales are equal.

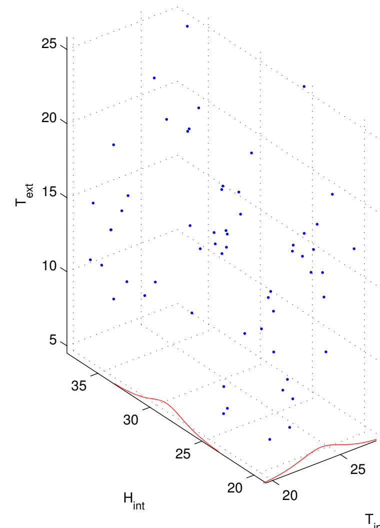

Figure 5. Data points following rescaling; note the axis scales are equal. Kernel shown in red. $(m_{t_{ext}} = 2.68, m_{h_{int}} = 0.62)$ **.**

heats up slowly. Statistically the deadband constitutes a *different process* which should be taken out of the data altogether. Figure 3 shows the same data as in Figure 2 but this time with the deadband removed. The relationship between internal temperature and Y is now somewhat clearer. Indeed a (toy) linear model (shown in Figure 3), created purely for illustration purposes (the actual model is a GP), shows a reasonable fit to the data. The toy model suggests that this users optimal internal temperature is in the region of 24°C, which is plausible. Given the preliminary elimination of the deadband data a more complex model, the GP, may now be constructed.

5.2 A Gaussian Process model.

Figure 4 shows the original *input* data for a single sample user: the internal temperature, the internal humidity and the external temperature. The original data shows that the internal temperature spans a lower range than the external temperature. [3] The Maximum Likelihood Estimates (MLE) for the scaling parameters (Equation 9) are $m_{T_{ext}} = 2.68$ and $m_{H_{ext}} = 0.62$.

Figure 5 shows the scaled input data for this user. It can see that the external temperature has in fact been expanded (by a factor of 2.68) while the Humidity has been compressed; overall the scales have changed quite significantly.

Figure 6 shows the estimated mean of the process over a grid of internal and external temperatures, i.e. $\hat{Y}_{x*}$ from Equation 11. There are a few things to note about this fit. First the estimated mean (Equation 11) drifts back to the process mean (Equation 13) as we get further from the sample points. This explains the counter-intuitive estimate which suggests the expected response to be zero at the boundaries.

[3] It is unclear how humidity is related to the temperatures as it is unitless (i.e. a percentage)

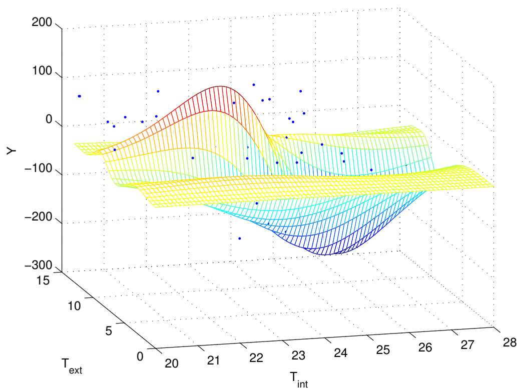

Figure 6. A fit of the preference data for user 1. (the humidity is kept constant at the average humidity, 36.2%.)

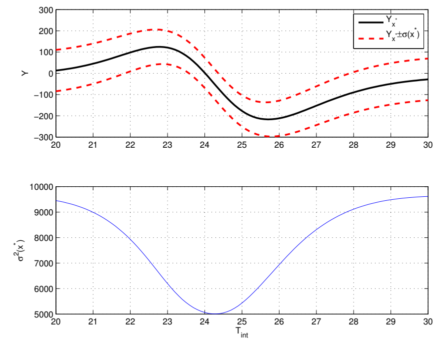

Figure 7. A fit of the preference data for user 1 with confidence intervals. (the humidity and external temperatures are kept constant at the average, 36.2% and 8.86°C, resp.).

For example, an internal temperature of 28°C is obviously too hot and so the mean estimate for Y should be approximatly -100 but as can be seen in Figure 6 the mean estimate is zero.

Figure 7 shows mean estimates generated from the same model used to produce Figure 6 but with the external temperature and internal humidity kept constant at their respective means. In addition, the variance at the evaluation points is also shown in the lower panel. As can be seen while the mean estimate, $\hat{Y}_x^*$ tends to zero as the internal temperature goes to 28° C, $\hat{\sigma}^2(x^*)$ increases from 5,000 to 10,000 reflecting the fact that $\hat{Y}_{x^*=28°C}$ has a high variance. Thus while $\hat{Y}_{x^*=28°C}$ is obviously wrong this is reflected clearly in the estimated variance at that point; $\hat{\sigma^2}_{x^*=28°C}$. The fit itself shows that this user is comfortable at 24°C but at 23°C the estimated response is 100 which lies between *slightly cool* and *cool*. The curve peaks at 23 °C as there is insufficient data at lower temperatures to make a good estimate (which is reflected in the higher variance).

5.2.1 Outlier influence reduction.

Initial estimates of the variance, ζ_x^0, are set according to $\zeta_x^0 = min(1/|Y_x|, 11)$. This equation reflects the fact that larger amplitude measurements are more likely to represent a deliberate response and was found to work quite well with this data set. Figure 8 illustrates three iterations of the outlier reduction algorithm. Figure 8 (a) shows the initial model fit. This user (different from that presented above) has been chosen as there is a large outlier, shown in red. As can be seen after three rounds the effect of this outlier has been significantly reduced.

5.3 Group results.

Table 2 shows parameter and hyperparameter estimates along with the number of samples for the UCC group. As the number of data points, n, increases, $\hat{\sigma}_x$ decreases; as expected. Note that this need not necessarily be the case; $\hat{\sigma}_x$ is not just a function of n but also the *quality* of the responses. For example, some users might return many responses but with large measurement errors. This would be reflected in a higher $\hat{\sigma}_x$. $\hat{\mu}_x$ is negative for all users indicating that the room is in general too hot for the group as a whole. The

Table 2. Statistics for the control group of users.

User	$\hat{\theta}$	$\hat{\nu}$	$\hat{m}_{T_{ext}}$	$\hat{m}_3$	$\hat{\sigma}_x$	$\hat{\mu}_x$	n
1	2.53	0.85	0.39	2.71	2.98	-9.43	61
2	4.97	2.79	0.13	2.01	14.60	-8.32	29
3	4.64	2.78	0.16	2.27	4.87	-19.43	48
4	5.00	2.26	0.11	2.18	29.69	-1.52	14

estimated hyperparameters for users 2 to 4 shows surprising consistency with $\hat{\theta} \approx 5$, $\hat{\nu} \approx 2.7$, $\hat{m}_2 \approx 0.13$ and $\hat{m}_3 \approx 2$.

5.3.1 Global priors

To examine the distribution of the hyperparameters the models from all users were collected together and the distribution of the hyperparameter estimates from *all* 65 users were examined to form *global* priors [3].[4] These are shown in Figure 9. While some outliers do exist (such as user 1 above) the distributions in Figure 9 are highly concentrated about their respective maxima. Specifically the maxima in Figure 9 are $\{\bar{\theta}, \bar{\nu}, \bar{m}_{T_{ext}}, \bar{\sigma}_x\} = \{4.91, 2.20, 0.63, 5.19\}$. This is a very useful result in several ways. First, there are physical interpretations for each parameter; $\bar{\theta}$ indicates that a poll taken at one point has an effective radius of 4.9° C (a useful rule of thumb for a grid survey). The scaling between internal and external temperatures is 0.63 and the variance of the 'preference process' is smooth at 5.19. Secondly, these global priors can be used as priors for a new user in order to give a reasonable fit before a significant amount of data has been collected.

5.3.2 Combining preference functions

Figure 10 shows the response curves for the 4 users in the UCC group. In addition the value of, $\hat{\sigma}^2(x^*)$, at the internal temperatures, x^*, is shown in the lower panel. The points at which $\hat{Y}_{x^*} = 0$ (the cross-over points) for the four users are $\{20.1, 22.9, 21.8, 21.0\}$, respectively. Thus the average cross-over point is $21.4°C$. However, this figure does not take into account that our confidence about each user differs significantly as can be seen on the lower panel of Figure 10. Figure 11 shows the probability distribution function of the response at 23°C. The data in this figure refers to just one

[4] Priors for $m_{H_{ext}}$ cannot be examined as there are no internal humidity readings for the CIT and NUIG sites.

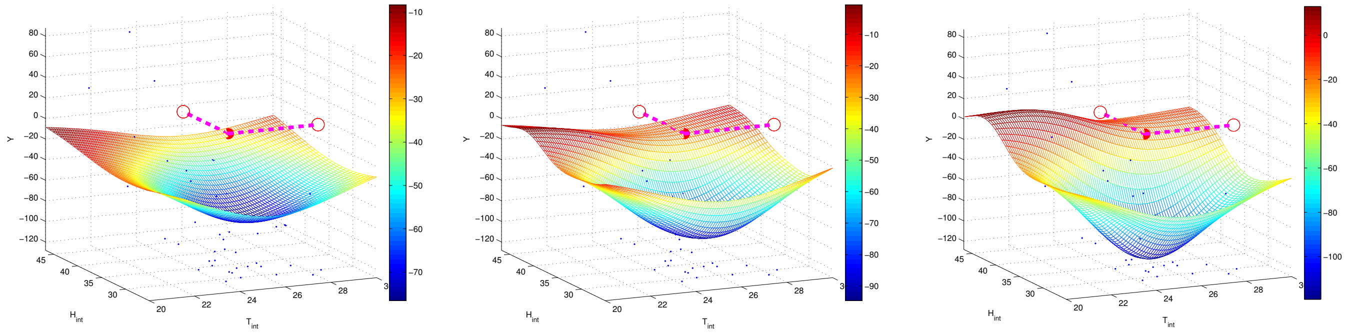

Figure 8. Outlier influence reduction;. a) The initial GP fit. b) after round 1. c) round 2.)

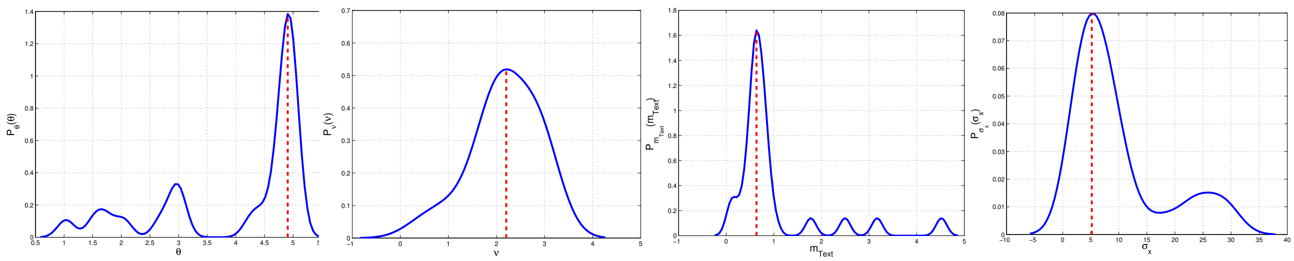

Figure 9. Global prior distributions for the hyperparameters.

operating point; the internal temperature at 23°C (the other variables being set to the average). Comparing Figures 10 and 11, it can be seen for example that User 4 has a mean response ($\hat{Y}_{x^*}$) of -10 at 23°C. In Figure 11 it can be seen that the distribution ($\sim \mathcal{N}\{\hat{Y}_{x^*}, \hat{\sigma}^2(x^*)\}$) around this response is quite narrow. For the other users the situation is different; although user 1 also has a response of $\hat{Y}_{23^\circ C} \approx -10$, $\hat{\sigma}^2_{23^\circ C}$ is much higher than for user 4 and so the corresponding distribution on Figure 11 is far wider. Here, we propose that the *expected number of neutral votes* as a combination function:

$$J_1(x^*) = \frac{1}{4}\sum_{i=1}^{4} P_i(Y_{x^*} = 0 | x = x^*) \qquad (16)$$

where $P_i(Y_{x^*} = 0 | x = x^*)$ is the pdf of the response for the i^{th} user evaluated at zero and $J_1(x^*)$ is the name given to this combination. Figure 12 shows the value of $J_1(x^*)$ corresponding to the response curves in Figure 10. It is satisfying to note that the peak occurs along a ridge and that the function decays smoothly away from the peak.

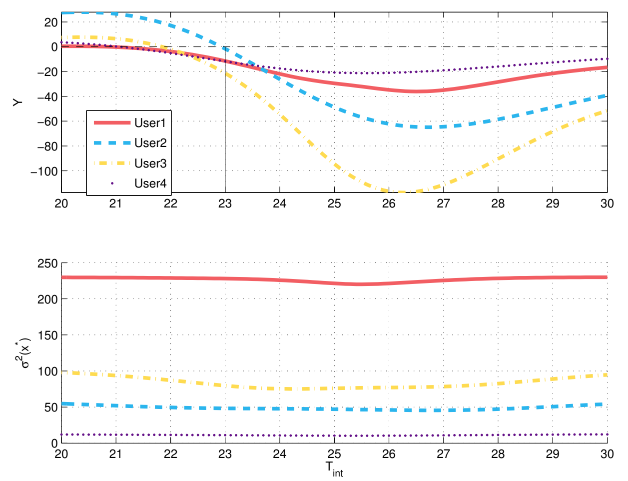

Figure 10. The response curve from the 4 control group users with respect to internal temperature. ($T_{ext} = 11°C$, $H_{int} = 39\%$)

6 Conclusions

For truly *smart* building operation, HVAC control should take user preferences into account. However, obtaining accurate comfort models tailored to specific users is difficult. We have developed a software app which requests users perceptions of the environment, and we show how to extract the meaningful data. We show how to extract individual preference functions from this noisy data using Gaussian Processes (GPs). Beliefs about the data can easily be integrated into GP models, from a prior belief about the population as a whole to belief about the measurement noise . We show how variance reduction can account for remaining outliers. The GP model also provides a *natural* means of combining the preferences of individuals sharing a space, to provide a usable signal for HVAC control. The focus here was on the data analysis. Future work will look at *how* and *when* the data should be collected. We are investigating an active learning approach, where $\hat{\sigma}^2(x^*)$ can be used to select an optimal time to poll the user. In addition, a GP model allows prior information to be easily integrated into the model (for example, there is no need to ask whether a user's response at 0°C is 'too cold'). Finally, we will investigate the application of GPs to other areas of intelligent building operation.

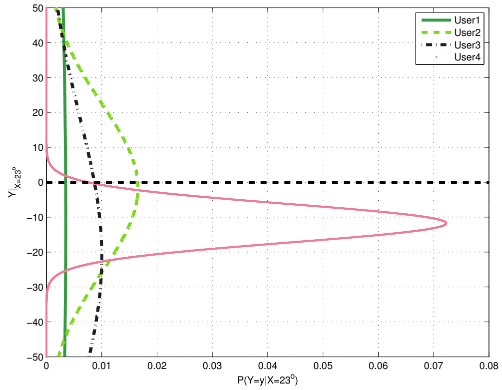

Figure 11. The distributions of the response at 23°C for the 4 users. ($T_{ext} = 11°C$, $H_{int} = 39\%$).

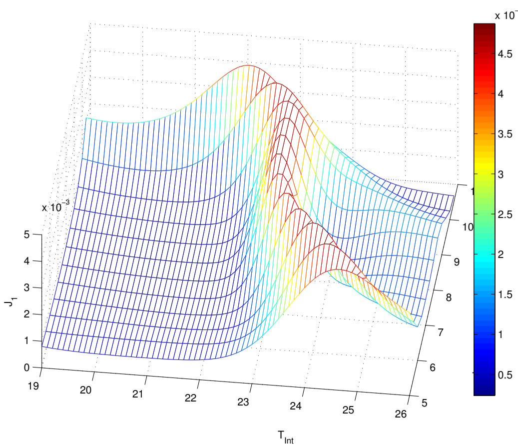

Figure 12. J_1 with respect to internal and external temperatures. (Humidity held constant at the average 38%.).

Acknowledgment

This work is funded by Science Foundation Ireland as part of the ITOBO project under grant No. 07.SRC.I1170.

References

[1] T. Alpcan. A framework for optimization under limited information. In *Proc. of the 5th Intl ICST Conference on Performance Evaluation Methodologies and Tools*, VALUETOOLS '11, pages 234–243, 2011.

[2] S. Atthajariyakul and T. Leephakpreeda. Neural computing thermal comfort index for hvac systems. *Energy Conversion and Management*, 46(15-16):2553–2565, 2005.

[3] R. E. Cass and D. Steffey. Approximate bayesian inference in conditionally independent hierarchical models (parametric empirical bayes models). *Jnl American Statistical Association*, 84(407):717–726, 1989.

[4] D. Daum, F. Haldi, and N. Morel. A personalized measure of thermal comfort for building controls. *Building and Environment*, 46:3–11, 2011.

[5] R. Dedear and G. S. Brager. Developing an Adaptive Model of Thermal Comfort and Preference. *ASHRAE Transactions*, 104(1), 1998.

[6] C. Duan, X. Ding, F. Shi, X. Xiao, and P. Duan. PMV-based fuzzy algorithms for controlling indoor temperature. In *IEEE Conference on Industrial Electronics and Applications*, pages 1492–1496, 2011.

[7] P. Fanger. *Thermal Comfort*. Danish Technical Press, 1970.

[8] P. W. Goldberg, C. K. I. Williams, and C. M. Bishop. Regression with input-dependent noise: A Gaussian process treatment. In *NIPS*, 1997.

[9] R. B. Gramacy. *Bayesian treed Gaussian process models, PhD thesis*. UC Santa Cruz, 2005.

[10] L. A. Jones and M. Berris. The psychophysics of temperature perception and thermal-interface design. In *Symposium on Haptic Interfaces for Virtual Environment and Teleoperator Systems'02*, pages 137–137, 2002.

[11] K. Kersting, C. Plagemann, P. Pfaff, and W. Burgard. Most likely heteroscedastic Gaussian process regression. In *Proceedings of the 24th Annual International Conference on Machine Learning (ICML 07)*, pages 393–400, Corvallis, OR, USA, June 20 2007.

[12] R. Martinez-Cantin, N. de Freitas, A. Doucet, and J. Castellanos. Active policy learning for robot planning and exploration under uncertainty. In *Proceedings of Robotics: Science and Systems*, Atlanta, GA, 2007.

[13] M. Orne. On the social psychology of the psychological experiment: With particular reference to demand characteristics and their implications. *American Psychologist*, 17:776–783, 1962.

[14] J. A. Orosa. A new modelling methodology to control hvac systems. *Expert Syst. Appl.*, 38(4):4505–4513, 2011.

[15] L. Prez-Lombard, J. Ortiz, and C. Pout. A review on buildings energy consumption information. *Energy and Buildings*, 40(3):394 – 398, 2008.

[16] C. E. Rasmussen and C. K. I. Williams. *Gaussian Processes for Machine Learning*. MIT press., 2006.

[17] T. J. Santner, B. Williams, and W. Notz. *The Design and Analysis of Computer Experiments*. Springer-Verlag, 2003.

[18] S. Seo, M. Wallat, T. Graepel, and K. Obermayer. Gaussian process regression: Active data selection and test point rejection. In *DAGM-Symposium*, pages 27–34, 2000.

[19] J. Serradilla, J. Q. Shi, and A. J. Morris. Fault detection based on gaussian process latent variable models. *Chemometrics and Intelligent Laboratory Systems*, July 2011.

[20] S. Wu and J.-Q. Sun. Two-stage regression model of thermal comfort in office buildings. *Building and Environment*, 57:88–96, Nov. 2012.