On Balancing Occupants' Comfort in Shared Spaces

Nic Wilson
Cork Constraint Computation Centre
Department of Computer Science,
University College Cork, Ireland
email: n.wilson@4c.ucc.ie

Abstract. Maintaining comfortable thermal conditions in an office environment is very important, as it can affect the quality of life of the occupants, their work productivity, and improve energy efficiency. One significant aspect of this task is how to balance the preferences of a number of occupants sharing the same space. We analyse and suggest some approaches to this problem, both for the case of optimising for a single time period, and for the problem of optimising over multiple different time periods.

1 INTRODUCTION

Maintaining comfortable thermal conditions in an office environment is clearly very important; apart from the quality of life of the occupants, it can also affect work productivity, and is very relevant for energy efficiency, since it can save considerable wastage through, for instance, overheating a room [15, 17]. One significant aspect of this task is how to balance the preferences of a number of occupants sharing the same space.

For a single occupant, one approach is, for each potential vector of environmental conditions, to produce a prediction of how they are likely to feel (e.g., a little warm), based on a scale for thermal sensation or comfort (see Section 2). We can then use this to generate a predicted dissatisfaction value. For example, if we are just using temperature to predict occupant responses, then each potential value of temperature is assigned a predicted dissatisfaction level, with a value of zero meaning completely comfortable. This therefore gives an objective function which we can aim to minimise. (This minimisation process can be complex, but is not the focus of the current paper.)

In Section 3 we consider the problem of evaluating overall occupant dissatisfaction for a single occupant over a number of (potentially temporally distant) time periods. Perhaps the main motivation for considering this situation is as a special case of the multiple occupants over time case. We argue that transforming the occupant dissatisfaction levels to a linear scale is very desirable.

We next (Section 4.1) consider the problem of evaluating overall occupant dissatisfaction for several occupants at a single timepoint (or time period). For multiple occupants, the notion of fairness comes into play. In Section 4.2 we consider the problem of evaluating overall occupant dissatisfaction for several occupants over a number of time-periods. The purpose of considering multiple timepoints is that it can be much easier to balance the occupant preferences over time, rather than for a single timepoint: we may be able to balance the mild discomfort felt by an occupant in one kind of scenario by what happens at other times. Section 5 concludes.

2 THERMAL COMFORT

Scales for Thermal Comfort: We want to predict how dissatisfied occupants are regarding thermal comfort in different environments. The standard ASHRAE seven-point thermal sensation scale [1] has integers -3 to $+3$ representing cold, cool, slightly cool, neutral, slightly warm, warm and hot, respectively. However, we are interested in occupant thermal comfort rather than thermal sensation, and they are not quite the same thing [2]. Because of this it is perhaps preferable to use labels that express thermal comfort more directly, such as: "Too cold", "slightly too cold", "a little cool", "comfortable", "a little warm", "too warm", "too hot". To allow expression of more extreme situations, we can add further points to the ends of the scale such as e.g., +4 for "much too hot", and +5 for "almost unbearably hot".

Predicting Thermal Comfort: Our approach relies on us being able to predict the thermal comfort of occupants based on different environmental conditions. Fanger's model [4] predicts the degree of a thermal sensation of an individual based on four physical variables and two personal variables: air temperature, air velocity, mean radiant temperature, relative humidity, clothing insulation and activity level. This mostly performs reasonably well [2, 7]. An alternative approach we have used is based on what the occupants have told us previously about how they felt in different situations. It appears that, at least in some cases, this can predict thermal sensation/comfort well, even if one only uses the air temperature data [14, 9].

3 OPTIMISING FOR A SINGLE OCCUPANT

Although our primary focus is multi-occupant rooms, it is helpful to first analyse the case of a single occupant.

3.1 Case of a Single Occupant for a Single Timepoint

Let us assume, at least for now, that we know, for any potential thermal conditions, the degree of thermal comfort for any occupant. The thermal conditions are measured by a collections of sensor readings, for example, including air temperature at different locations, humidity and so on.

Let $\vec{\theta}$ represent a vector of sensor readings. We are assuming we have a real-valued function H on all such vectors $\vec{\theta}$ (or at least on all feasible ones), that correctly predicts the degree of thermal comfort $H(\vec{\theta})$ for the occupant in situation $\vec{\theta}$. Thus if, for example, $H(\vec{\theta}) = 1$

then it is predicting that the occupant will be a little warm in the situation represented by $\vec{\theta}$.

Clearly, what we'd like to do is to control the system so as to generate conditions $\vec{\theta}$ that make $H(\vec{\theta})$ as close to 0 (fully comfortable) as possible.[1] However, we need to decide more precisely what it means for $H(\vec{\theta})$ to be close to 0. We want some function L which maps values of thermal comfort to degrees of dissatisfaction. For example, if we'd like the function to be symmetric about 0 (the y-axis) (i.e., an even function), we might choose $L(x) = x^2$ or $L(x) = |x|$, in which cases we will aim to optimise $H(\vec{\theta})^2$ or $|H(\vec{\theta})|$, respectively. In contrast, some occupants may dislike more being cold than being warm, leading to a function L that is not symmetric about the y-axis.

3.2 Case of Single Occupant over Time

It can be important to consider a longer term view than just optimising for a single timepoint (or time period, where we assume that each time period is the same length). We consider here the case of evaluating a collection of comfort scores (or, alternatively, dissatisfaction scores) for a single occupant at different timepoints. For example, suppose a particular control policy leads to thermal comfort/sensation levels scores 3 at 9am, 0 at 2pm and 0 at 4pm. A different control policy leads to 1 at 9am, 1 at 2pm and 1 at 4pm. Which of these two is better? They each have the same average comfort value, so at first sight one might consider them to be equivalent. However, the score of 3 may indicate the occupant being much too hot, possibly being very uncomfortable. Because of this, one can argue that a score of 3 is really much worse than three times as bad as a score of 1, with the latter indicating at most mild discomfort.

More generally, suppose that the occupant registers ASHRAE-scale (or thermal comfort) values over a number $x_1, \ldots, x_N$ of timepoints; we want to evaluate these to give an overall cost value. We want some function L that takes these sequence of thermal comfort values and generates a dissatisfaction score, which is on an additive scale—i.e., the combination of a collection of dissatisfaction scores is their sum. Thus the overall evaluation (dissatisfaction) of a sequence of comfort values $x_1, \ldots, x_N$ will then be defined to be $\sum_{i=1}^{N} L(x_i)$, so that for two sequences $(x_1, \ldots, x_N)$ and $(y_1, \ldots, y_N)$ if $x_1 + \cdots + x_N = y_1 + \cdots + y_N$ then the two sequences are equally adequate. We call L, a "linearising function".

We may have $L(-x_i) = L(x_i)$ but we might not. Also, the cost (negative utility) of, for example, "almost unbearably hot" can be very different for different occupants in different situations; in particular, it makes a major difference if the occupant can acceptably leave the room and find somewhere more comfortable.

3.2.1 Generating Linearising Function L

Standard elicitation procedures for utility [13] can be adapted to elicit values of the function L. For example we could ask queries such as: *For which value of x is* $(0, 0, x)$ *equivalent to* $1, 1, 1$*?* This could be used to specify $L^{-1}(3)$, given that we've defined $L(0)$ to be 0 and $L(1)$ to be 1.

Note that positive linear transformations of the scores make no essential difference, so that if L is an adequate linearising function, then so is $DL + E$, given by $(DL + E)(x) = DL(x) + E(x)$, where D and E are strictly positive real numbers. This allows us to normalise in some way the function L if we wish, by transforming L using such a linear transformation in order to ensure certain conditions are respected. Firstly, we can normalise L to ensure $L(0) = 0$. We could also, if we wished normalise to ensure that e.g., $L(1) = 1$; or e.g., $L(1.5) = 1$.

Candidate Functions

If we do not the have the opportunity to elicit values for L, a standard function can be used. A simple candidate function L is given by $L(x) = x^2$, with e.g., $L(2) = 4$, so a comfort value of 2 is judged to be four times as bad as a comfort value of 1. One possible family of functions, which can bias much more to the more extreme points of the scale, is that of the form: $L(x) = (a^{x^2} - 1)/(a - 1)$ for some $a \geq 0$ with $a \neq 1$, where we set $L(1) = 1$. L is continuous and $L(0) = 0$.For example, with $a = 1.2$, $L(2)$ is around 5.4 and $L(3) \approx 20.8$. For a close to 1, $L(x)$ tends to x^2. For $a > 1$, $L(x)$ grows faster than x^2; for $a < 1$, $L(x)$ grows slower than x^2.

3.2.2 Further Interpretation of Linear Scale

In summary, we should attempt to map user discomfort values to a linear scale. A further advantage of using a linear scale is when we are summing up previous data (or uncertainty about discomfort values) with an expected value. The latter doesn't make so much sense if the scale is not a linear one.

The transformed values, on the linear discomfort scale, can be considered as negative utility (in the sense of expected utility). One interpretation of such values is as the financial gain they would require in order to suffer this degree of discomfort (this relates somewhat to "Willingness to Pay", Section 3.8 of [13]). This financial interpretation might seem a little far fetched at first sight: there is unlikely to be any differential compensation paid to occupants according to their degrees of comfort. However, thinking about a home situation, the relevant kinds of comparisons are being made implicitly. If I am working at home on my own on a cold winter's day I will tend to adjust the heating so that I am slightly but not very cold. I am thus implicitly expressing a tradeoff between my comfort level and the cost of fuel: I am paying for a rate of fuel usage that transforms my thermal comfort level from too cold to slightly cool; but I'm implicitly not being prepared to pay for the extra amount of fuel to move my thermal comfort level to completely comfortable.

4 MULTIPLE OCCUPANTS

In this section, we are focusing on multiple occupants who share the same space, such as in the same room, or perhaps in adjoining rooms with thin partitions separating the rooms. This does not mean that the different occupants are experiencing the same thermal conditions; for example, the air temperature around one occupant may be 22^oC, but 20^oC for another occupant in the same large office, for example, if the former is receiving direct sunlight.[2] (Our approaches would apply also for multiple occupants in different rooms that are far apart; however, the method is not so relevant there, since for the latter situation the control system could treat the two occupants independently, with an action relevant for one of the occupants, such as turning on a heater in her room, not relevant for the other.)

[1] We may well also want to involve energy efficiency or energy usage in the objective function, which can be done by adding a separate term. We do not focus further on this issue in this paper.

[2] Collaborators on the ITOBO project [8] are currently performing experiments within two multi-occupant rooms in the Environmental Research Institute building at University College Cork, where various parameters including air temperature and lux (light) level are measured at different points in the room.

For each occupant in each situation let us again assume that we know what their dissatisfaction level is (we consider the case of where there is uncertainty briefly later in Section 4.4). We also assume that these dissatisfaction levels are on a linear scale, as described in Section 3.

4.1 Single Decision For Multiple Occupants

First we consider a situation where we are interested in optimising at a single timepoint (time period), for multiple occupants sharing the same space.

We would like to treat the occupants fairly. Firstly, we would like to regard them each as having equal importance (although later, in Section 4.3, we consider allowing different grades of importance). Secondly, we would like to bias towards having a more equitable range of degrees of discomfort. For example, in a two occupant office, we would prefer a situation where both occupants have degrees of dissatisfaction of 1 (on a linear scale) than one having degree of dissatisfaction of 0, and the other degree of dissatisfaction of 2.

4.1.1 Relative Scaling of Occupants for Equal Occupant Importance

The linearising functions L described in Section 3 were, for the purpose of optimising for a single occupant, non-unique, in that one can multiply the function L by a strictly positive real scalar, and get a function that performs equivalently. For multiple occupants, we will have such a scaling function L_i for each occupant i. We are considering a situation where the occupants are assumed to be of equal importance. We therefore need to scale the different functions L_i for each occupant so that they reflect equal importance.

We are assuming, as described above, that for any vector $\vec{\theta}$, representing the environmental conditions, a degree of discomfort/dissatisfaction $J_i(\vec{\theta})$, which is a non-negative real number. In particular, J_i is based on a function L_i that maps a thermal comfort level to a non-negative real number. In order to respect the requirement that the occupants have equal importance, we need a way of calibrating the functions L_i for different occupants. As mentioned above, we can ensure that $L_i(0) = 0$, i.e., that thermally neutral (comfortable) corresponds with a zero degree of dissatisfaction. We want to rescale the functions L_i to reflect equal importance.

One approach is to normalise the functions L_i in some way; for instance to ensure $L_i(1) = 1$; or alternatively, $L_i(2) = 1$; or $L_i(1.5) = 1$.

Alternatively, if the functions L_i can be given a financial interpretation, as the monetary value needed to compensate for the thermal discomfort, then this already gives a relative calibration/scaling of the functions L_i.

One might perhaps also try to take into account, in this relative scaling step, the relative “choosiness” of different occupants: some occupants will tend to give more extreme inputs, which will tend to have more impact on the objective function. One could decide to correct for this, effectively lessening the importance of such occupants, so that their inputs will tend to dominate less.

After the rescaling process, for each vector of environmental conditions we have a vector $(s_1, \ldots, s_K)$ of dissatisfaction values, one for each of the K occupants in the currently considered space.

4.1.2 Properties of Aggregation Operators

Our task is to sum the vector $(s_1, \ldots, s_K)$ of dissatisfaction values (one for each occupant) into a single non-negative real number that somehow represents the overall degree of (thermal) dissatisfaction for the group. We thus want to define some real-valued function G that sums up these K numbers into a real value $G(s_1, \ldots, s_K)$. Such an aggregation function G can be a kind of average or a generalised summation.

The topic of aggregation operators and their properties has been studied for a long time, at least since Cauchy in 1821 and Kolmogoroff and Nagumo in 1930 [12, 10]. We give some properties that are arguably natural for our problem.

The first property implies that the result is not affected by the identity of the occupant. It also implies that the occupants are equally important.

Symmetric: if $(t_1, \ldots, t_K)$ is a permutation of $(s_1, \ldots, s_K)$ then $G(t_1, \ldots, t_K) = G(s_1, \ldots, s_K)$.

The next two properties are very natural: if one increases the degree of dissatisfaction of any user then the overall dissatisfaction is increased (or at least not decreased, for the former property).

Increasing: $G(s_1, \ldots, s_K)$ is an increasing function of its arguments, i.e., if for all $i = 1, \ldots, K$, $s_i \leq t_i$ then $G(s_1, \ldots, s_K) \leq G(t_1, \ldots, t_K)$.

Strictly Increasing: $G(s_1, \ldots, s_K)$ is a strictly increasing function of its arguments, i.e., if for all $i = 1, \ldots, K$, $s_i \leq t_i$ and for some i, $s_i < t_i$, then $G(s_1, \ldots, s_K) < G(t_1, \ldots, t_K)$.

A small increase in input dissatisfactions should not cause a large jump (continuity). The second property is less clearly essential, but reflects the smoothness of the change in output as the input changes.

Continuous: G is a continuous function;

Continuously Differentiable: G is continuously differentiable

The following property from [6] represents the idea that if we bring the values s_i closer together without changing their sum then we improve the overall evaluation.

Transfer Principle: If $s_i < s_j$ and $\epsilon < s_j - s_i$ then $G(s_1, \ldots, s_K) > G(t_1, \ldots, t_K)$ where $t_i = s_i + \epsilon$, and $t_j = s_j - \epsilon$, and $t_k = s_k$ for $k \neq i, j$.

The following properties relate to the fact that the choice of the linearising functions L_i (see Section 3.2.1) would often be non-unique, in that a positive linear transformation of the same function could also do. The second is perhaps less important, since we are arranging that $L_i(0) = 0$. The two properties together are known as being *stable for positive linear transformations* [12].

Scaling: For $C > 0$,

$$G(Cs_1, \ldots, Cs_K) = C \times G(s_1, \ldots, s_K).$$

Uniform translation: For $D > 0$,

$$G(s_1 + D, \ldots, s_K + D) = G(s_1, \ldots, s_K) + D.$$

The following seems natural for an averaging operator:

Idempotence: $G(s, \ldots, s) = s$ [only for an averaging operator].

The next property relates to sum-like operators, and is a convenient associativity property:

Decomposable as binary operator: [only for sum operators] G can be decomposed as an associative binary operation $\oplus$: there exists associative binary operation $\oplus$ such that $G(s_1, \ldots, s_K) = s_1 \oplus \cdots \oplus s_K$.

Ordinary Summation and Mean

We can define the sum operator $G(s_1, \ldots, s_K) = s_1 + \cdots + s_K$, or the mean operator $G(s_1, \ldots, s_K) = \frac{1}{K}(s_1+\cdots+s_K)$. This satisfies all the above properties except the Transfer Principle. However, for the sake of fairness, we want to ensure that the Transfer Principle is satisfied.

We go on the suggest three families of aggregation operators that might be used in this context, and briefly discuss their properties.

4.1.3 Ordered Weighted Average

One idea is to use a weighted average, where more weight is attributed to give extra weight to the more uncomfortable occupants. This kind of average is known as an ordering weighted averaging operator [16, 5]. If there are K occupants present then one uses positive decreasing weights $w_1, \ldots, w_K$ that sum to 1. The overall cost function is then equal to $w_1 s_{(1)} + \cdots + w_K s_{(K)}$, where $(s_{(1)}, \ldots, s_{(K)})$ is a permutation of $s_1, \ldots, s_K$ such that $s_{(1)} \geq \cdots \geq s_{(K)}$.

Since there may not be constant number of occupants in the same space, we need a sequence[3] $w_1, \ldots, w_K$ of such weights for each K.

We can generate a corresponding Sum operator, by multiplying through by K:

$$G(s_1, \ldots, s_K) = Kw_1 s_{(1)} + \cdots + Kw_K s_{(K)}$$

If we set $w_1 = 1$ and otherwise $w_i = 0$ we obtain the max operator. Also if we set w_i proportional to ϵ^i for some small positive number ϵ we get an operator that orders vectors in a similar way to the leximin ordering [3]. However, both of these give what might be considered as excessive weight to the most uncomfortable occupant.

For any weights vector, the corresponding aggregation operators are symmetric, increasing and continuous, and satisfy both Scaling and Uniform Translation. Idempotence is also satisfied for the averaging version. The operators are strictly increasing when all the weights are non-zero, and satisfies the Transfer Principle whenever all the weights are different.

It is typically not decomposable as a binary operator. The only other one of the above properties not satisfied by ordered weighted averages is being continuously differentiable. The derivative with respect to s_i at points when $s_i \neq s_j$ for all $j \neq i$, is one of the weights (specifically w_j such that $s_i = s_{(j)}$). The derivative is thus discontinuous (except if all the weights are equal).

4.1.4 Skewing of Linear Scale

Let λ be a continuously differentiable (strictly) increasing bijection on the non-negative reals.

We can define a Sum operator G_λ by

$$G_\lambda(s_1, \ldots, s_K) = \lambda^{-1}\Big(\sum_{i=1}^{K} \lambda(s_i)\Big).$$

An averaging operator can be defined similarly:

$$\lambda^{-1}\Big(\frac{1}{K}\sum_{i=1}^{K} \lambda(s_i)\Big).$$

This is known as a quasi-arithmetic mean [11].

[3] For different K, one would expect the sequences to be related in some way, so one might consider coherence conditions that relate to the weights as K varies; however, we do not consider this issue further in the current paper.

We are interested especially in cases where λ has a strictly increasing derivative, i.e., a positive second derivative (except possibly at 0) since then it satisfies the Transfer Principle. The smoothness properties are satisfied given suitably smooth λ (e.g., strictly positive derivative, except possibly at 0). G_λ doesn't generally satisfy Scaling and Uniform Translation. However, if we use $\lambda(x) = x^a$ for some $a > 1$ then Scaling is satisfied. The other properties are satisfied.

4.1.5 Mean-Plus-Spread Approaches

Another approach is to consider the value of vector $(s_1, \ldots, s_K)$ as consisting of two components: the first being the mean $\frac{1}{K}(s_1 + \cdots + s_K)$, and the second relating to the spread of the s_i's around the mean. For example, we can consider functions G of the form

$$G(s_1, \ldots, s_K) = \mu + R\sigma,$$

where R is a non-negative real number (which may depend on the number of occupants K), $\mu = \frac{1}{K}(s_1+\cdots+s_K)$ is the mean of the K values, and σ is their standard deviation, so that $\sigma^2 = \frac{1}{K}\sum_{i=1}^{K}(s_i - \mu)^2$, which equals $\frac{1}{K}\sum_{i=1}^{K}(s_i)^2 - \mu^2$.

G is strictly increasing if we choose R such that $R < 1/\sqrt{K-1}$. It is not decomposable as a binary operator, but satisfies the other properties.

4.2 Multiple Occupants With Multiple Time Periods

We now consider a situation where we have multiple occupants in the shared space, for multiple time periods. We would like to generate an objective function that gives an overall cost (overall degree of dissatisfaction/undesirability).

We assume that the thermal sensation input for each occupant has been mapped to a linear scale of dissatisfaction, as in Section 3. We can therefore sum these degrees of dissatisfaction to get an overall degree of dissatisfaction for each occupant. Different occupants may be present different numbers of time periods, and it can be natural sometimes to explicitly take this into account. and it is important to take this into account, For each occupant we therefore then have a pair (s_i, N_i) representing the summed degree of dissatisfaction s_i and the number N_i of time periods in which they were present. Define s_i^* to be s_i/N_i, the mean degree of dissatisfaction for occupant i.

We therefore would like to generate a function F that takes as input a sequence

$$\Big((s_1, N_1), \ldots, (s_K, N_K)\Big),$$

representing the summarised inputs for the K occupants.

The three families of aggregation operators of Section 4.1 can be adapted for this task. A simple way to optimise for the overall function is, at each timepoint, to control for the environmental conditions so as to minimise $F\Big((s'_1, N'_1), \ldots, (s'_K, N'_K)\Big)$, where $\Big((s'_1, N'_1), \ldots, (s'_K, N'_K)\Big)$ corresponds to the inputs received so far. On the other hand, if we were to have information about the expected occupancy and dissatisfaction scores in future periods, then the optimisation technique could take these into account.

4.2.1 Ordered Weighted Sum Approach

The idea here is again to use a weighted sum/average, where higher weights are attributed to more dissatisfied occupants. So, we choose decreasing weights $w_1, \ldots, w_K$. We define the overall cost function F by

$$F\Big((s_1, N_1), \ldots, (s_K, N_K)\Big) = Kw_1 s_{(1)} + \cdots + Kw_K s_{(K)},$$

where $(s_{(1)}, \ldots, s_{(K)})$ is a permutation of $s_1, \ldots, s_K$ such that $s^*_{(1)} \geq \cdots \geq s^*_{(K)}$. Note that the weight assigned to the ith occupant is based on their mean degree of dissatisfaction, $s^*_i = s_i/N_i$, but the overall cost is based on their total degree of dissatisfaction s_i.

4.2.2 Skewing of Linear Scale

The approach from Section 4.1.4 can be applied directly. Again let λ be a function with the properties given in Section 4.1.4 (e.g., continuously differentiable strictly increasing bijection on the non-negative reals, which has a strictly increasing derivative). We define the overall cost function F_λ by

$$F_\lambda\Big((s_1, N_1), \ldots, (s_K, N_K)\Big) = \lambda^{-1}\Big(\sum_{i=1}^{K} \lambda(s_i)\Big).$$

Here the values N_i do not come into the definition.

As before, we can use, for example, λ of the form $\lambda(x) = x^a$ where $a > 1$.

4.2.3 Mean-Plus-Spread Approaches

The approach described in Section 4.1.5 can be adapted easily. We can consider a random variable which, for $i = 1, \ldots, K$, takes value s^*_i with chance $\frac{N_i}{N}$ where $N = \sum_{i=1}^{K} N_i$. Let μ be the mean of this random variable and σ^2 be the variance, so that

$$\mu = \sum_{i=1}^{K} \frac{N_i}{N} s^*_i = \frac{1}{N} \sum_{i=1}^{K} s_i,$$

and

$$\sigma^2 = \sum_{i=1}^{K} \frac{N_i}{N} (s^*_i - \mu)^2.$$

Again we can define the overall cost function to be $\mu + R\sigma$, choosing positive real R (which may depend on N).

4.3 Incorporating Importance

The approaches in Section 4.2 can easily be adapted to take an importance weight $v_i > 0$ into account for each occupant i. Each time period that occupant i spends in the space is treated as v_i time periods to give more emphasis to occupants with higher importance weight. We replace s_i by $v_i s_i$ and N_i by $v_i N_i$, and apply the equations in Sections 4.2.1 4.2.2 and 4.2.3 to obtain objective functions that bias according to the importance weights.

4.4 Taking Uncertainty into Account

There are many potential sources of uncertainty in our application. For instance, uncertainty about:

- what the current environmental conditions are, because of inaccuracy in sensing;
- the thermal sensation value the occupant will feel in any given conditions; this is true if we use a learning algorithm based on past inputs, or a PMV-based approach;
- each particular occupant's mapping from the thermal sensation scale to degrees of dissatisfaction;
- which occupants will be present.

One common and natural approach to dealing with this uncertainty is to use some kind of expected value, specifically for dissatisfaction given particular environmental conditions. This gives another reason to use a linear scale for dissatisfaction; if we compute expected value on a non-linear scale then the value will not necessarily adequately sum up the distribution.

Alternatively, we can generate probability distributions over the dissatisfaction levels of each occupant, giving a random variable for each occupant's dissatisfaction. The different methods described above for computing the values of an objective (cost) function can be extended for this probabilistic case, using, for example, a Monte-Carlo algorithm to estimate expected cost, if we assume the occupants' random variables are mutually independent.

5 CONCLUSION

The paper addresses the issue of how one evaluates a collection of thermal comfort inputs from multiple occupants and over time. This is important for defining an objective function for control, based on predicted responses of occupant under various environmental conditions.

We argue that it is important firstly to map the degrees of thermal comfort onto a linear scale of dissatisfaction, so that the values can be summed for the combination of several values for a single occupant. We have suggested three families of approaches for aggregating the dissatisfaction scores of several occupants, with different strengths and weaknesses. We have shown how this may be applied for the case of multiple occupants over many time periods. Although all three families seem quite natural, there are, of course, other approaches that should be explored.

There are other issues that could be considered for more sophisticated approaches. For instance, we assumed that the ordering of the sequential inputs was unimportant; however, for consecutive time-points this could make a difference to the overall dissatisfaction of an occupant, where, for example, starting off cold and slowly increasing heat would presumably be better than a more random ordering.

ACKNOWLEDGEMENTS

This material is based upon works supported by the Science Foundation Ireland under Grant No. 08/PI/I1912, and the ITOBO Strategic Research Cluster. I am grateful to the reviewers for their comments, and for discussions with many collaborators on the ITOBO project including Ken Brown, Anika Schumann, Conor Ryan, Damien Fay, Massimo Manna and Mateo Burillo.

REFERENCES

[1] ANSI/ASHRAE Standard 55-2004, 'Thermal environmental conditions of human occupancy', (2004).

[2] K.E. Charles, 'Fanger's thermal comfort and draught models', Technical Report IRC-RR-162, National Research Council of Canada: Institute for Research in Construction, (2003).

[3] D. Dubois, H. Fargier, and H. Prade, 'Refinements of the maximin approach to decision-making in fuzzy environment', *Fuzzy Sets and Systems*, **81**, pp. 103–122, (1996).

[4] P.O. Fanger, *Thermal Comfort: Analysis and Applications in Environmental Engineering*, McGraw-Hill, NY, 1972.

[5] J. Fodor, J.-L. Marichal, and M. Roubens, 'Characterization of the ordered weighted averaging operators', *IEEE Transactions on Fuzzy Systems*, **3**(2), 236–240, (1995).

[6] Christophe Gonzales, Patrice Perny, and Jean-Philippe Dubus, 'Decision making with multiple objectives using gai networks', *Artif. Intell.*, **175**(7-8), 1153–1179, (2011).

[7] M. A. Humphreys and J. F. Nicol, 'The validity of ISO-PMV for predicting comfort votes in every-day thermal environments', *Energy and Buildings*, **34**(6), 667–684, (2002).

[8] ITOBO, 'Information & communication technology for sustainable and optimised building operation (http://zuse.ucc.ie/itobo/)', *Cork, Ireland*, (2007-2012).

[9] C. Manna, N. Wilson, and K. Brown, 'Learning individual thermal comfort by robust locally weighted regression with adaptive bandwidth', in *ECAI-2012 Workshop: AI Problems and Approaches for Intelligent Environments*, (2012).

[10] J.-L. Marichal, *Aggregation Operators For Multicriteria Decision Aid*, Ph.D. dissertation, University of Liège Liège, Belgium, 1998/1999.

[11] J.-L. Marichal, 'On an axiomatization of the quasi-arithmetic mean values without the symmetry axiom', *Aequationes Mathematicae*, **59**(1-2), 74–83, (2000).

[12] J.-L. Marichal, P. Mathonet, and E. Tousset, 'Characterization of some aggregation functions stable for positive linear transformations', *Fuzzy Sets and Systems*, **102**(2), 293–314, (1999).

[13] R.L. Keeney H. Raifa, *Decisions with Multiple Objectives: Preferences & Value Tradeoffs*, Cambridge University Press, 1993.

[14] A. Schumann, M. Burillo, and N. Wilson, 'Predicting the desired thermal comfort conditions for shared offices', in *Proceedings of the International Conference on Computing in Civil and Building Engineering (ICCCBE-10)*, pp. 95–96, (2010).

[15] A. Wagner, E. Gossauer, C. Moosmann, Th. Gropp, and R. Leonhart, 'Thermal comfort and workplace occupant satisfaction–results of field studies in german low energy office buildings', *Energy and Buildings*, **39**(7), 758–769, (2007).

[16] Ronald R. Yager, 'On ordered weighted averaging aggregation operators in multicriteria decisionmaking', *IEEE Trans. Syst. Man Cybern.*, **18**(1), 183–190, (January 1988).

[17] R. Yao, B. Li, and J. Liu, 'The impact of office environments on employee performance: The design of the workplace as a strategy for productivity enhancement', *Building and Environment*, **44**, 2089–2096, (2009).