Joint International Conference on Computing and Decision Making in Civil and Building Engineering
June 14-16, 2006 - Montréal, Canada

UTILITY BASED DECISION MAKING IN BUILDING INFRASTRUCTURE

Vipul Singhvi[1], James Garrett[2], Scott Matthews[3], Guestrin[4], Khee Poh Lam[5]

ABSTRACT

Meeting human preferences of comfort, safety and privacy are major factors in success of many infrastructural projects. The inability to integrate the seemingly important factor of human preference in decision making is due to lack of a framework to quantitatively capture these preferences. In this paper, we present a utility-theory based approach to successfully integrate these preferences in decision making. Utility theory is a micro-economic concept used to measure the happiness or satisfaction gained from a good or service. We use this same concept and define a utility function to summarize the individual preferences for comfort, safety and privacy. We address the challenges of defining a realistic utility function, optimizing the resulting decision theoretic problem and integrate that into a formal decision making approach for building operation.

KEYWORDS

Utility theory, human preferences, optimization, lighting control, decision making

INTRODUCTION

Meeting human preferences of comfort, safety and privacy are major factors in the success of many civil-infrastructural projects. Currently, these factors are taken into account by incorporating available standards during decision making. In building operation, occupant's comfort is measured using standards like ASHRAE (ASHRAE 1980) and other such agencies. However, most of the standards represent approximation of these preferences, as in reality, they are unique for each individual and often are location and time dependent. The inability to integrate the seemingly important factor of human comfort in building operation is due to lack of a framework to quantitatively capture their comfort preferences. In this paper, we present a principled decision theoretic approach using utility theory to successfully integrate these preferences in building operation. Utility theory (Varian 1992) is a micro-economic concept used to measure the happiness or satisfaction gained from a good or service. The principled decision theoretic approach is used to implement an intelligent lighting control that elicits occupant's preferences and then optimally controls the lighting system.

[1] Graduate Student, Dept. of Civil Engineering, Carnegie Mellon University, 5000 Forbes, Pittsburgh, Phone: +1 (412) 268 2086, FAX: +1 (412) 268 7813, vipul@cmu.edu

[2] Professors, Dept. of Civil Engineering, Carnegie Mellon University, 5000 Forbes, Pittsburgh, Phone: +1 (412) 268 5674, FAX: +1 (412) 268 7813, garrett@cmu.edu

[3] Assistant Professor, Dept. of Civil Engineering, Carnegie Mellon University, 5000 Forbes, Pittsburgh, Phone: +1 (412) 268 2914, FAX: +1 (412) 268 7813, hsm@andrew.cmu.edu

[4] Assistant Professor, School of Computer Science, Carnegie Mellon University, 5000 Forbes, Pittsburgh Phone: +1 (412) 268 3075, FAX: +1 (412) 268 3075, guestring@cs.cmu.edu

[5] Professor, School of Architecture, Carnegie Mellon University, 5000 Forbes, Pittsburgh, Phone: +1 (412) 268 8503, FAX: +1 (412) 268 7819, kplam@cmu.edu

Building operation is a complex activity, where building operators and occupants continuously interact with each other. Typically, the interaction is passive with little communication between the operator and the occupants. Historically, *maximizing occupant comfort* and *minimizing energy costs* have always been two primary objectives of intelligent buildings operation (Finley Jr. M. R. 1991; Flax B. 1991). The trade-off between meeting occupant preferences for indoor environmental condition and reduction in energy usage leads to a difficult optimization problem and this optimization can be thought of as *decision-theoretic* — the objective is to *minimize* the expected cost of building operation and to *maximize* the occupant's expected comfort. Utility theory is a principled decision theoretic approach that can be used for such optimization problems (Keeney 1976). Utility theory stems from the notion of *preferences* over *outcomes*. Outcomes result from the choices made by the system or the occupant. The occupant's comfort preferences are defined in terms of an order, f , over these outcomes. This preference order can be defined in terms of a real-valued *utility function*, **U**, over outcomes — one prefers state $\mathbf{s_1}$ over $\mathbf{s_2}$ (written $\mathbf{s_1}$ f $\mathbf{s_2}$) if and only if $\mathbf{s_1}$ has higher comfort and hence utility: $\mathbf{U}(\mathbf{s_1}) > \mathbf{U}(\mathbf{s_2})$. The advantage of this optimization-based approach is the potential to *personalize* the system for individual occupants, simply by defining separate utility functions for each user. The requirement of minimizing operation cost can be modeled as a utility function which would monotonically decrease with the increase in the operation cost of the system.

In this paper, we focus on a utility based formal decision making for lighting in commercial buildings. US office buildings use over 86 billion kWh for lighting each year (CBECS 2003). At the average energy cost of $0.08 per kWh, the potential savings from implementing energy efficient lighting in 50% of office buildings is more than $2.1 billion per year. However, at the same time, energy efficient lighting is typically associated with reduced lighting which can negatively affect the productivity of the occupants. Reduced productivity and costs incurred due to loss of work can significantly outweigh the benefits from saving energy. Singhvi et al. (Singhvi 2005) present a decision-theoretic approach to optimally achieve occupants' light preferences and energy usage tradeoff by solving a multi-criterion optimization problem. They show that given the utility function for individual occupants and the operation cost their *coordinated illumination approach* can efficiently optimize the tradeoff in meeting occupant preferences and energy usage.

While decision-theoretic optimization provides a powerful, flexible, and principled approach for such systems, the quality of the resulting solution is completely dependent on the accuracy of the underlying utility function. Unfortunately, defining a good utility function, a process called *utility elicitation*, is a complex, time consuming, and an error-prone task. Utility elicitation is a critical process for the success of a decision theoretic system. In this paper, we present a utility elicitation approach that addresses the needs and requirements of an intelligent building system. In addition, we extend the approach presented by Singhvi et al. by defining a more realistic utility function for lighting comfort. The main contributions of this work are:

- A formal decision theoretic framework for integrating occupants' preferences in building operations
- A principled utility elicitation technique using minimal interaction and partial information provided by the occupants; and
- A novel interface design for recording occupants' qualitative preference

UTILITY THEORY: LIGHT CONTROL

Utility functions are defined over a space which is exponential in the number of variables on which the utility depends. Typically such representation of utility function leads to intractable optimization problem due to exponential nature of the solution space. More tractable representations of the utility are possible if we make certain assumptions about *additive independence* (Keeney 1976) among the variables. These assumptions allow the function to be decomposed into smaller components, thus reducing the number of parameters needed to specify it completely. In this paper, we assume that the occupant's utility function is derived from *sub-utility* components, which reflects preference of the occupant for various parameters in the indoor environment. More formally we assume that the occupant's utility function **U** is linearly additive, i.e. there is a set of sub utility function $\boldsymbol{\varphi} = \{\phi_1, \phi_2 ....\phi_k\} \in [0,1]^k$**,** such that for any indoor environment state **s** in the building system the associated utility for the occupants is given by:

$$\Phi(\mathbf{s}) = \sum_{i=1}^{n} w_i * \phi_i(s_i) \qquad (1)$$

Here, $\mathbf{s}$ is the indoor environment state defined as the vector of the various parameters s_i involved in defining the state. $\Phi(\mathbf{s})$ is the occupant's utility function representing the preference for the comfort in state **s.** ϕ_i is the sub-utility function and w_i are the associated weights in the utility function.

In the formulation presented in Singhvi et al. each occupant, '*k*', has a utility function $\Phi_k(\mathbf{s})$ representing their preference for a given light setting **s**. In their definition of **s**, they have used a restrictive definition of lighting state by using only one parameter, the horizontal light intensity (Eq. 2). Figure 1 shows typical utility function used for the operation cost and occupant preference.

$$\Phi_k(\mathbf{s}) = w_1 * \phi_1(Lux) \text{ where } w_1 = 1 \qquad (2)$$

The *operating cost utility function* is defined as Ψ , which decreases monotonically with the energy expended for maintaining the state **s**.

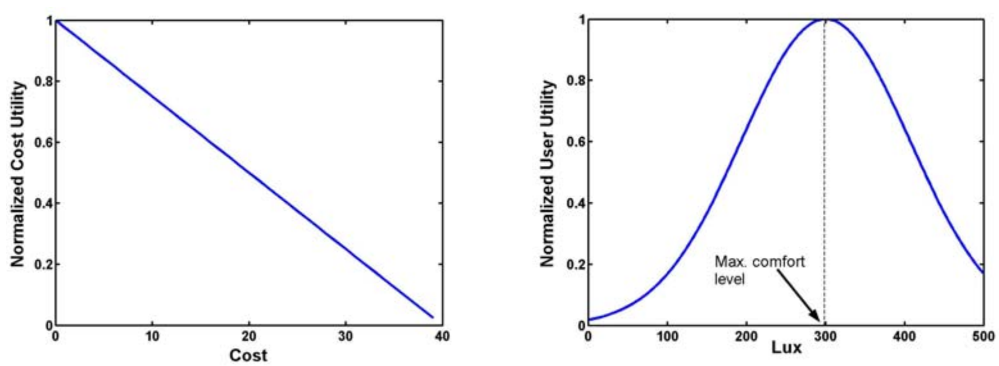

Figure 1: Typical utility function (Singhvi et al., 2005)

A building has multiple occupants with varying preferences and hence with varying utility functions $\Phi_1, \Phi_2......\Phi_n$. When considering occupant preferences and the operating cost, the goal is to tradeoff Ψ with the occupant utility $\sum_{i=1}^{m} \Phi_i$. They use a common technique of '*scalarization*' to solve this multi-criterion optimization problem by defining a *system utility* function, $U(\mathbf{s})$ which is defined as follows:

$$U(\mathbf{s}) = \sum_{i=1}^{m} \Phi_i(\mathbf{s}) + \gamma * \Psi(\mathbf{s}) \qquad (3)$$

The optimal control strategy in this scenario to find **s*** such that:

$$\mathbf{s}^* = \underset{\mathbf{s}}{\operatorname{argmax}}\, U(\mathbf{s})$$

Note that the state **s** is exponential and enumeration of all possible states is impossible. Singhvi et al. present an efficient algorithm to solve this exponential maximization problem exploiting the *zoning* principle in lighting design.

The success of such decision support system depends on how well the utility function Φ represents the preferences of the occupants. The main challenge is to estimate/elicit the shape of the *sub-utility* function ϕ_i (ref: eq 1) and the associated weights w_i for each occupants. These factors are unique for individual occupants and need to be estimated to develop a reliable utility model for the occupants. In the rest of the paper, we start off by discussing related research in the area of utility elicitation. We then identify the major requirements in a utility elicitation tool for building operations and present our design of the utility elicitation tool.

PREFERENCE / UTILITY ELICITATION

Previous work in the area of user preference elicitation and tradeoff analysis has generally followed two main approaches, *classical and behavioral decision theory* (Pu 2003). The classical theory deals with the idea of formulating a perfect model of users' preference utility function. Given a perfect utility function, the classical theory can accurately predict final configuration of a configurable item or service. Classical decision theory (Keeney 1976) treats tradeoff problems under the assumption that a machine is able to help a human to externalize a value structure and use it to evaluate decision outcomes. A popular method to elicit such value function is to ask users to choose a set of outcomes and infer the model from their choices. This process in general can be lengthy and cognitively demanding. One such method uses *gamble queries*, asking the user whether they prefer choice '*X*' to a mixture of a probability p chance of the ideal outcome and a (1 − p) chance of the worst outcome. While this and similar techniques (Chajewska 2000; Chajewska 2001) have elegant theoretical properties, we feel occupants cannot coherently report their preferences with respect to probability distributions over lighting requirements and these processes would be very lengthy.

Behavioral decision theory (Payne 1993; Carenini 2002), on the other hand, is very concerned with decision makers' behavior. Many years of studies have pointed out the adaptive and constructive nature of human decision making. Although individuals clearly aim at maximizing the accuracy of their decisions; they are often willing to tradeoff accuracy to reduce cognitive effort. Stating preferences is a process rather than a one-time enumeration of preferences that do not change over time and also user involved preference construction is likely to be more effective than using default or implicit models if a user is to understand and accept the solution outcomes (Carenini 2002). Several decision support systems have been using behavioral approach to integrate user's preferences in different domains. Systems like FindMe(R. Burke 1997), ATA (Linden 1997), and AptDecision (Shearin 2001) help users to navigate through a large space of alternatives to find the most preferred solution.

We seek a utility elicitation tool that differs significantly from the tools discussed above in terms of *design assumptions*, *implementation* and *usability* requirements. The utility elicitation tools discussed above assume an implicit sub-utility function or existence of expert knowledge which reduces the task to estimating the associated weights to form the full utility function. Most of the time the assumptions made about the shapes of sub-utility functions are valid, for example, monotonically decreasing utility function with respect to cost of a product. However, due to unique personal preferences of occupants these assumptions cannot be used in building operation domain. We feel for the occupants to trust the building operation system, the system should let the occupants define their *sub-utility* function.

In terms of implementation and usability requirements, the utility elicitation tools presented above are designed where users are engaged in an active dialogue with the system. The users interact with the system with a tangible goal in mind, for example, choosing a flight, finding an apartment and so on and the system has their full attention during the interaction period. On the contrary, occupant's interaction with the building controls is at best passive. Even though occupants spend most of their time in the building, by nature they seldom interact actively with the building system. Designing a successful tool for building environment requires additional requirements in terms interaction functionality. Due to restrictive interaction between occupants and building controls, we feel a successful building control system should:

- be cognitively less demanding on occupants
- require minimal interaction time
- be able to operate using an approximate utility function
- rely on observing occupants behavior to progressively fine tune the utility function; and
- be flexible to the changing requirements of occupants

In the design of the utility elicitation tool, we feel occupants will likely report more accurate preferences in a setting where their context or the state is visible. While occupants are capable of specifying preferences between concrete outcomes, they have difficulty articulating a real-valued utility function. We seek to design a tool where occupants can convey their qualitative preferences based on real visible context. The tool would provide a quantitative interpretation of their preferences to the system. We feel the quantitative interpretation of the user's qualitative feedback should be visible so that occupants can trust and relate to the system's interpretation of their requirements. This is important for the system to be accepted as this keeps the control in hands of the occupants. The design should free the occupants from having to reason about numerous and unintuitive parameters, probabilities or monetary values of different tradeoffs and thus makes it cognitively easier for the occupants.

To address the unique requirements and features required for a utility elicitation tool, we have designed our tool in two stages. In the first stage called, ***sub-utility elicitation***, we let the occupant visually explore the lighting space by letting them operate the lighting control, while constantly soliciting qualitative feedback. Since every occupant has a unique sub-utility function, it is important that each occupant provide the system with that information. This is the only stage where we solicit active interaction from the occupants. The goal of this stage is to estimate the shape for all the sub-utility functions.

The second stage is called ***weight elicitation***. In this stage, the system starts by using a set of weights to form the utility function for the occupants. The weights are chosen to meet certain constraint based on the domain knowledge and decision theoretic concepts. In this stage, the occupant does not have to interact with the utility elicitation system, however if an occupant is unsatisfied he/she can use the lighting control to change the setting. The utility elicitation system observes the behavior and uses it to update the utility function. Over time with minimal interaction from the occupant the system would be able to learn a very close approximation of the utility function. In the next section we describe the two stages of utility elicitation in more detail.

STAGE 1: SUB-UTILITY ELICITATION

For defining the utility function for the lighting comfort of occupants we use following three parameters of indoor environments:

1. Horizontal light intensity (Lux)
2. Glare
3. Uniformity

Using the assumption of additive linearity (Keeney 1976) we formally define the utility function for occupant comfort as follows:

$$\Phi(\mathbf{s}) = w_1 * \phi_1(Lux) + w_2 * \phi_2(Glare) + w_3 * \phi_3(Uniformity) \qquad (4)$$

Here, w_i is the weight factor, which signifies how much the i^{th} factor influences the overall utility function, where $\phi_{1,2,3} \in [0,1]$. The goal of this stage is to estimate the shape of the sub-utility function for each of the three parameters.

The sub-utility elicitation tool can automatically construct a good approximation of the sub-utility function by soliciting feedback about the lighting state. Since the preference for the indoor environments are ephemeral the system is flexible to adapt to the changing requirements by letting the occupants update their sub-utility function at any time. The real time interaction provided in the interface makes it cognitively easier for the occupants to identify the effect of the parameters in the sub-utility function. We feel this is a better way of eliciting the sub-utility function as opposed to asking them queries to make monetary tradeoffs for various lighting scenarios.

The sub-utility interface contains three main modules: *preference module*, *control module* and the *state module*. The preference module contains two axes, y-axis representing the preference level and the x-axis represents the measure for the three parameters (light intensity, glare or uniformity respectively). The control module provides the occupants with the control to increase or decrease the corresponding parameters. The state module provides information about the current lighting intensity, glare and uniformity index.

The goal of the interface is to provide real time information about the current lighting state while the occupant is exploring the parameter space available in the given workspace. The occupant starts exploring the parameter space by using the control module. For every lighting state that he/she chooses to evaluate, the interface provides him with real time information about the lighting state and solicits his preference through the preference module. The system requires the occupant to evaluate at least 4-5 states. Based on the conveyed qualitative preference, the system presents a quantitative interpretation of the occupant's utility function by superimposing the interpretation on the preferences entered by the occupant in the preference module. We assume that the occupants' utility function can be closely approximated by using a third degree polynomial.

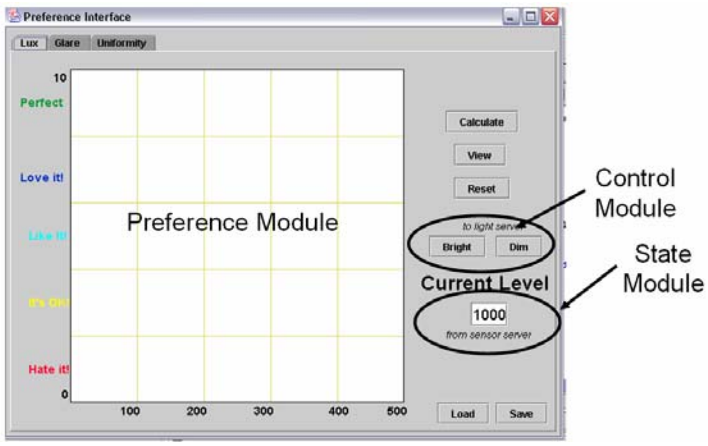

Figure 2: Sub Utility Elicitation Interface

Schematically we can view the process in two phases. The first is called *visual exploration* and the second automated phase is called *sub-utility estimation phase*. The following figure shows example of two different occupants interacting with the sub-utility interface.

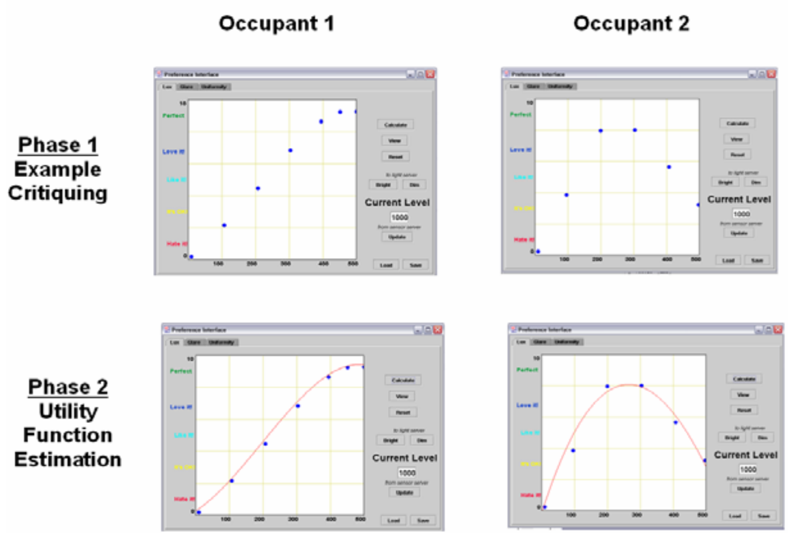

Figure 3: Interface for interactive sub-utility elicitation

We normalize the utility functions obtained from this stage to ensure that each of the sub-utility function obtained meet the condition $\phi_{1,2,3} \in [0,1]$.

STAGE 2: WEIGHT ELICITATION

Stage 1, sub-utility estimation, provides with the model for the sub-utility function for the three parameters in the utility equation. However, in equation 4 we need estimate the weights to form the full utility function for the occupants. Estimating even a close approximation of these weights would take significant time of the occupants. At the same time, the success of such a system is based on how well the utility function represents the occupant's preferences. Here, we use an approach to learn the best estimate of the utility weights by observing the occupant's interaction with the system. The lighting system starts of by assigning random set of weights for the utility function. The random weights chosen must satisfy the following condition: are based on some domain knowledge and

$$\sum_{i=1}^{n} w_i = 1 \tag{5}$$

Since the system is designed to provide lighting in a workspace, we assume the preference of the occupant more for lighting level (w_1) than for glare (w_2) and since glare affects the user more then uniformity we assume $w_2 > w_3$. However, these are not hard constraints and the occupant can change it anytime later on. The lighting system uses the weights generated using these constraints. These weights do not reflect the true preferences of the occupants however they provide a good starting point for the lighting system. The system is designed to update the weights by observing the behavior of the occupants. The updating mechanism is based on the choices the occupants makes while trying to changing the lighting condition. Whenever the occupant changes the lighting state set by the main system, we use this information as his *preference ordering* over the two states. For example, if the occupant changes the current state $\mathbf{s}_1$ to say state $\mathbf{s}_2$, we interpret his choice in form utility function as follows:

$$\Phi(\mathbf{s}_2) > \Phi(\mathbf{s}_1)$$

$$\left(\sum_{i=1}^{n} w_i * \phi_i\right)_2 - \left(\sum_{i=1}^{n} w_i * \phi_i\right)_1 > 0$$

We define a choice vector **c** such that,

$$\mathbf{c}_{ij} = (\phi_k)_j - (\phi_k)_i \;\; for\; k = 1,...,n \;\; and\; i \neq j$$

so, we have

$$\mathbf{w}.\mathbf{c}_{ij} > 0 \tag{6}$$

We refer to equation 6 as *contextual information*. Occupants can convey other types of contextual information, which the system can utilize while calculating the best utility estimate based on the current context. We consider three types of partial context information (Malakooti 2000):

1. **(type 1)** Lower and upper bounds, LB_i and UB_i for each attribute weight factor w_i can be provided by the occupant.
$$LB_i \leq w_i \leq UB_i \;\; for\; i = 1,...,n$$
2. **(type 2)** Ranking of pairs of weight factors is provided by the occupants:
$$w_i > w_j \;\; for\; i, j = 1,...,n \; i \neq j$$
3. **(type 3)** Paired comparison of some alternative states provided by the occupants, that is if state $\mathbf{s_i}$ is preferred to $\mathbf{s_j}$ then
$$\phi(s_i) - \phi(s_j) > 0$$
$$\mathbf{w}.\mathbf{c}_{ij} > 0$$

We represent the *contextual information* as a set $\mathbf{\Lambda}$, which forms a constraint set on **w**. The system tries to estimate the value of weights which satisky the constraint set. To calculate the estimate of these weights we use the concept of *utility non-domination* (Malakooti 2000)

An alternative $\mathbf{s}$ *is utility dominated if and only if there exist* $\mathbf{s}'$ *such that* $\Phi(\mathbf{s}) < \Phi(\mathbf{s}')$ *for all* $\mathbf{w} \in \Lambda$ *.Otherwise,* $\mathbf{s}$ *is utility non-dominated.*

Figure 4: Utility domination definition

To get the set of weights for the given contextual information we define the linear programming approach as follows:

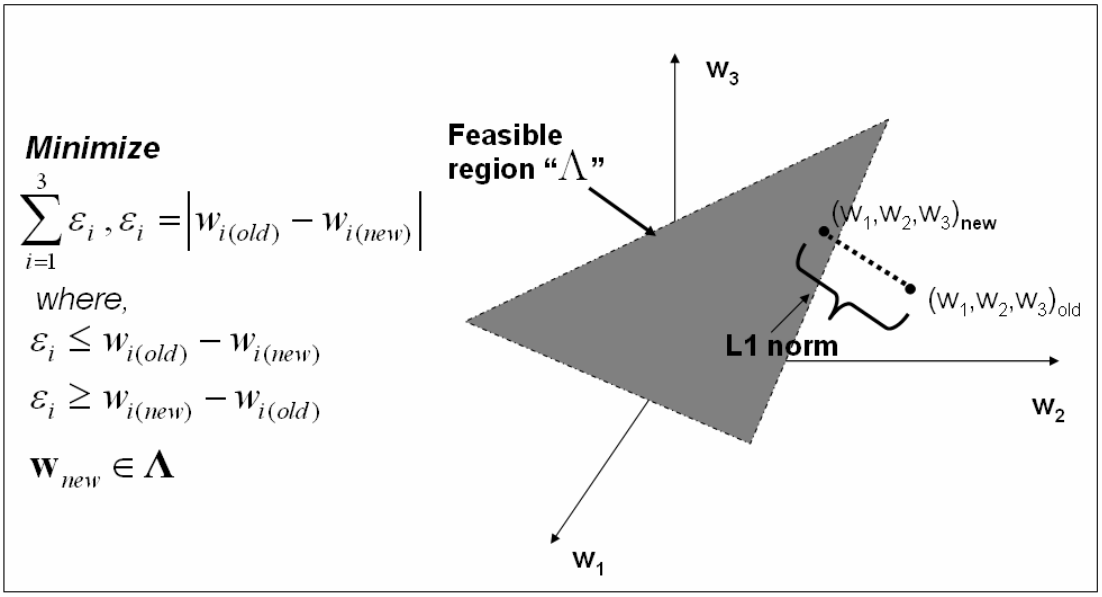

Figure 5: Convex optimization formulation for weight calculation

The approach tries to minimize the L1 norm defined as $\sum_{i=1}^{3} \varepsilon_i , \varepsilon_i = \left| w_{i(old)} - w_{i(new)} \right|$. If this LP problem is infeasible, then the original assumption of linearity in the utility model is violated. Otherwise, any solution will correspond to weight estimate that meet the occupant's preferences. This weight update technique is a conservative estimate as it tries to find the set of new weights that is closest to the current set of weights. We adopt the solution **w** as the new weight estimate. We assume that the occupant's sub-utility function remains constant, so whenever the occupant convey a choice that is inconsistent with the current sub-utility function the system would ask the occupant to update the utility function or asks the permission to disregard the choice that violates the utility function. To test the preference elicitation approach we are currently controlling two real life work spaces to do detailed user studies.

CONCLUSION

In this paper, we presented a principled approach to integrate occupants' preferences into formal decision making for building operation. The main goals in building operation are meeting occupant comfort and reducing operating cost. Simultaneously meeting these goals leads to a complex optimization problem. Singhvi et al. (Singhvi 2005) present a utility based decision-theoretic optimization approach to solve the problem. While decision-theoretic optimization provides a powerful approach for such systems, the quality of the resulting solution is completely dependent on the accuracy of the underlying utility function.

In this paper, we presented a utility elicitation technique that can be used to elicit occupants' preferences for indoor lighting. We formulate utility elicitation technique into a two step process, sub-utility elicitation and weight elicitation. In sub-utility elicitation the occupants can visually explore the lighting environment while providing qualitative preferences for various lighting states. The system utilizes the qualitative information to create quantitative interpretation of the preferences. In the second stage the system uses the partial information in terms of constraints on the weight to identify a feasible region for allowable weights. It then uses a complex optimization formulation to identify a set of feasible weights that meet the occupant's preferences. We present an implementation of the approach in form of an intelligent light control. The utility elicitation process requires minimal interaction and is cognitively less demanding on occupants. The system is able to operate on partial information provided by the occupants and adepts by observing the actions of the occupants.

REFERENCES

ASHRAE (1980). American Society of Heating Refrigeration and Air Conditioning Engineers, Atlanta.

Carenini, G., Poole, D. (2002). Constructed Preferences and Value-focused Thinking: Implications for AI research on Preference Elicitation. . AAAI - 02 Workshop on Preferences in AI and CP: symbolic approaches, Edmonton, Canada.

CBECS (2003). Commercial Building Energy Consumption Survey. D. o. E. Energy Information Administration.

Chajewska, U., Koller, D., (2000). Utilities as random variables: Density estimation and structure discovery. UAI.

Chajewska, U., Koller, D., Ormoneit, D., (2001). Learning an agent's utility function by observing behavior. ICML.

Finley Jr. M. R., A. K., and R. Nbogni (1991). "Survey of intelligent building concepts." IEEE Communication Magazine.

Flax B., M. (1991). "Intelligent buildings." IEEE Communication Magazine.

Keeney, R., L., Raiffa, H. (1976). Decisions with Multiple Objectives: Preferences and Value Trade-offs. New Yorl, Wiley.

Linden, G., Hanks, S., Lesh, N. (1997). Interactive assessment of user preference models: The automated travel assistant. In Proceedings of User Modeling.

Malakooti, B. (2000). "Ranking and Screening Multiple Criteria Alternatives with Partial Information and use of Ordinal and Cardinal Strength of Preferences." IEEE Trans. on Systems, Man, and Cybernetics Part A **30**: 355-369.

Payne, J., W., Bettman, J.,R., Johnson, E.,J. (1993). The Adaptive Decision Maker, Cambridge University Press.

Pu, P., Kumar, P., Faltings B., (2003). User-Involved Tradeoff Analysis in Configuration Tasks. Third International Workshop on User-Interaction in Constraint Satisfaction, Ninth International Conference on Principles and Practice of Constraint Programming.

R. Burke, K. H., and B. Young (1997). "The findme approach to assisted browsing." IEEE Expert **12(4)**: 32–40.

Shearin, S., Lieberman, H. (2001). Intelligent Profiling by Example. Conference on Intelligent User Interfaces, ACM Press.

Singhvi, V., Krause A., Guestrin, C., Garrett, J., Matthews, S., (2005). Intelligent Light Control using Sensor Networks. SenSys, San Deigo.

Varian, H. R. (1992). Microeconomic Analysis.