

2022-02-04

A psychological model for the prediction of energyrelevant behaviours in buildings: Cognitive parameter optimisation

von Grabe, J

<http://hdl.handle.net/10026.1/18758>

10.1049/ccs2.12042

Cognitive Computation and Systems

Wiley Open Access

All content in PEARL is protected by copyright law. Author manuscripts are made available in accordance with publisher policies. Please cite only the published version using the details provided on the item record or document. In the absence of an open licence (e.g. Creative Commons), permissions for further reuse of content should be sought from the publisher or author.

A Psychological Model for the Prediction of Energy-relevant Behaviors in Buildings: Cognitive Parameter Optimization

Jörn von Grabe¹, Sepideh Korsavi²

¹Institute for Architecture and Planning, University of Liechtenstein, Vaduz, Liechtenstein

²Department of Built Environment, University of Plymouth, Plymouth, Devon, UK

Emails: v.grabe@buildingsimulation.eu (J.G.), sepideh.korsavi@plymouth.ac.uk (S.K.)

Abstract

Energy consumption in buildings is a major contributor to global warming and therefore has become a field of intensive research. This type of energy consumption can be described in two dimensions: an appliance-based dimension and a behavior-based dimension. To address the behavior-based dimension a recent study proposed a cognitive human-building interaction model that builds on the instance-based learning paradigm. However, since the values of the standard cognitive parameters commonly used for modelling lab-based behaviors are not suitable for the “real-world” domain of human-building interaction, this paper aims to identify cognitive parameter values adapted to and suitable for the specific character of this application domain.

To achieve this goal, a virtual test environment - consisting of an occupied room and a corresponding model task - was designed to test the performance of the model and its dependence on a set of fundamental cognitive parameters. A test criterion was developed that did not depend on empirical data but used the predictive consistency of the model as reference. A range of values was pre-selected for each parameter based on theoretical and empirical considerations, which was then tested against the evaluation criterion.

The performance of the model was improved significantly throughout the parametrization process and yielded plausible results.

Keywords: Cognitive Modelling; Energy-relevant Behavior; Instance-based Learning; Prediction

| Nomenclature | | | | | |
|--------------|-----|--|---------|-----|--|
| Acronyms | AQ | Air quality | Symbols | A | Total activation |
| | BV | Blended Value | | B | Base level activation |
| | jnd | Just noticeable difference | | d | Decay factor |
| | PMV | Predicted Mean Vote | | e | Past encounters of chunks |
| | RF | Relaxation factor | | fan | Number of associations between source slot and slots in memory |
| | PM | Penalization of activation due to partial matching | | H | Heating state |
| | pm | Partial matching scaling parameter | | k | Weber constant |
| | SA | Spreading activation | | N | Noise activation |
| | SD | Standard deviation | | n | Number of chunks in memory |
| | SG | Sub-goal | | P | Probability of retrieval of an instance |
| | SIM | Similarity parameter for slot comparison | | R | Result stored to an instance |

| | | | | | |
|---------------|--------|---------------------------------|--|----------|--|
| | BPS | Building performance simulation | | r | Total number of instances (belonging to an action) |
| | ABM | agent-based modelling | | s | Total number of slots of an instance |
| | PMV | Predicted Mean Vote | | S | Associative strength (w/o index: maximum associative strength) |
| | TRNSYS | Transient Systems Simulation | | T | Thermal |
| | RF | Relaxation Factor | | | |
| Abbreviations | avg | Average | | t | Time |
| | CAT | Category | | U | Utility of an instance |
| | clo | Clothing value | | x | New sensation experience |
| | Dim | dimension | | γ | Random draw out of [0,1] |
| | met | Metabolic rate | | σ | Noise scaling parameter |
| | clo | Clothing | | τ | Imprecision of retrieval (temperature parameter) |
| | curr | Current | | i, j | i_{th} and j_{th} instance/ slot |
| | heat | Heating | | m | m_{th} action |
| | Sm | Stimulus | | p | p_{th} occurrence of an instance in the past |
| | wind | Window | | | |
| | Sim | Similarity | | | |

1. Introduction

The energy consumption of a building includes both technical dimensions related to its construction and appliances and human dimensions related to occupants' interaction with the construction and the appliances. Consequently, attempts to reduce the energy consumption of buildings and their contribution to greenhouse gas emissions should consider occupants' energy-relevant behaviors.

Building performance simulation (BPS) tools are thermodynamic-based numerical simulation programs used to virtually test and optimize the energy performance and indoor environmental conditions of a building. To consider the behavior dimension of energy consumption, numerical models for behaviors such as adjusting thermostats, switching lights or opening and closing windows are developed and integrated into BPS tools.

Most approaches for developing numerical energy-related behavioral models use uni- or multivariate probabilistic models (e.g., Bernoulli-, survival- or Markov models) established based on a statistical analysis of behavioral and environmental observations. For example, observed window operation is statistically correlated with indoor and outdoor temperatures, blind control with solar intensity or domestic lighting demand with internal brightness (see [1-18] or [19, 20] for a more detailed analysis of such studies). Through co-simulation with building and environmental models, such statistical behavior models predict the probability of the state or state change of a control at a certain time under certain environmental conditions. As an alternative to such macromodels, agent-based modelling (ABM) approaches represent single occupants as "agents" who autonomously interact with their environment. These agents' behaviors are usually based on a priori defined behavior hierarchies or algorithmic rules that can also be derived from empirical data. For example, empirical records of fan and heating usage and corresponding environmental conditions are converted into behavior rules representing a hierarchy of actions under specific conditions which is then used for simulation (for examples, see [21-23]). Both types of approaches have in common, first, that they are correlative

models that do not use psychological methods, and second, that they are specific to the context in which the behavior was observed. The latter implies that they are static in their rules and lack the ability to adapt to new conditions like real human behavior. Despite the plethora of such frameworks that researchers have developed since the 1980s-1990s, a general and valid interaction model has not yet emerged [24-27].

A recently suggested alternative to currently established methods is the system-theoretical approach to modelling energy-relevant behaviors [28], which shares some similarities with ABM approaches in its intent to model individual behaviors while it is different from probabilistic techniques in adopting a psychological perspective. The system-theoretical approach to modelling action provides a framework and a reference system that permits a detailed analysis of all the psychological processes involved in constituting an action such as attention, perception, deciding, problem-solving or learning [29]. Most importantly, the proposed model facilitates psychological learning regarding action consequences in a particular environment, which in turn provides knowledge that informs the execution of future actions. This provides models with human-like flexibility that enhances behavioral adaptation to a variable environment.

This study focuses on the action planning, decision and learning processes involved in human-building interaction. These processes were implemented as an instance-based learning and decision process and recently described in detail in [28]. However, the algorithms used to compute the involved cognitive processes include a set of cognitive parameters [30-32], the value of which had to be estimated to suffice the specific characteristics of a model that is not applied to a short laboratory experiment but to a real-world task that spans over thousands of hours of simulated time.

Since the validity and applicability of the cognitive model is highly dependent on the appropriateness of the cognitive parameters, they must be selected carefully. This paper suggests and illustrates a substantiated process through which appropriate parameters can be identified. It starts with the introduction of the methodology, which in Section 2 firstly explains the model structure and secondly introduces the procedure of the parameter study. Section 3 presents the essential part of the study, i.e., the cognitive parameter study, which is followed by a summary and conclusion in Section 4.

Please be aware that, even though the suggested model has potential to be utilized in comfort-based building automation systems, the suggested process is not to be confused with the PCAO algorithm (Parametrized Cognitive Adaptive Optimization). The PCAO algorithm is generally used to optimize complex large-scale problems and has also been suggested to be used in energy and comfort management in automated buildings [33].

2. Methodology

To understand the process of optimizing the cognitive parameters of the cognitive model, it is first important to understand the general set up of the simulation model. This includes the structure of the action model (the human model that perceives, decides and acts) as well as its connection to the environmental model (which represents both, the physical structure of the space and the variable thermodynamic and comfort conditions). A schematic representation of this set-up is illustrated in Fig.1. In this figure, the cognitive model, for which the parameters are to be optimized, is represented by the “planning & decision” sub-process.

The necessary details of the action model are explained in Section 2.1 whereas the relevant details of the environmental model are illustrated in Section 2.2. Section 2.3 explains the procedure of the cognitive parameter study. Since it is of utmost importance to clearly understand this procedure, this section is further subdivided into the scope of the parameters to be optimized (Section 2.3.1) and the criteria used to evaluate the appropriateness of the parameter values (Section 2.3.2).

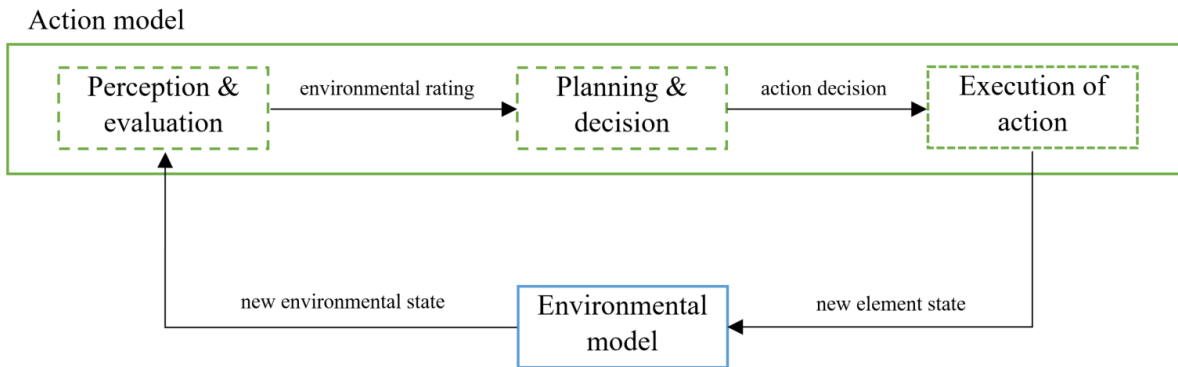


Fig 1: Schematic diagram of the action model [28].

2.1 Action model

A psychological model of goal-directed action can be understood as a process connecting at least three distinct psychological sub-processes, all of which need to be numerically modeled in order to obtain an adequate action model: 1. the perception and evaluation of the action-relevant environmental conditions; 2. the planning of and decision for action (the cognitive model); and 3. the execution of the action. This sequence of sub-processes is briefly described in the following paragraphs.

Sensation models are used to predict the perception and evaluation of the action-related environment. Though multiple energy-relevant needs have been identified [34], the current model considers only two needs: thermal sensation and indoor air quality (IAQ) sensation, which are both part of the perceivable context for the model. This simplification ensures that the behavior of the model is comprehensible and interpretable while still sufficiently complex to simulate occupant's decisions.

The thermal sensation model is based on Fanger's predicted mean vote (PMV) model [35] that approximates the rating of the thermal environment on the ASHRAE thermal sensation scale, a 7-point scale ranging from -3 (representing cold sensations) across 0 (representing neutral sensations) to +3 (representing hot sensations). The PMV algorithm is based on the thermodynamic quantities of air temperature, air humidity, radiant temperature and air velocity, among which the first three variables are dynamically computed by the environmental model (Section 2.2) and the fourth is estimated to be 0.1 m/s. Additionally, the PMV algorithm considers the metabolic rate of the occupants—estimated to be 1.2 met to correspond to sedentary office activity—and their clothing, which is measured in the unit clo and adjusted by the model based on its own decisions.

The IAQ sensation model is an adaptation of Fanger's self-decipro concept [36], as modified through later findings by Gunnarsen and Fanger [37]. This adaptation yields the air quality sensation on a scale comparable to the ASHRAE thermal sensation scale; however, it has a narrower range from 0 (representing ideal conditions) to 3 (representing very low air quality).

The sub-process of planning and deciding, the cognitive model, is implemented as an instance-based learning model [30, 31] that has been adapted and extended to enable its application to the complexities of human-building interaction. This paradigm assumes that information about context-action-result contingencies is stored as instances to the declarative memory during action. This information can later be retrieved in comparable contexts to inform appropriate actions in light of the current goal. The various submodels and algorithms used to set up the entire decision model are too complex to be summarized here. Some of the ruling equations will be referred to throughout this paper; however, the reader is referred to [28, 38] for a more detailed explanation.

Finally, the execution of the action itself is reduced to a “just do” model; once the decision model has chosen an action, the determined action state is immediately set. This simplification means, for example, that the movements of occupants through space to operate the building elements or the detailed action of adjusting the clothing is not explicitly modelled.

2.2 Environmental model

The environmental response to the actions of the model is determined by a thermodynamic building model, implemented in TRNSYS (transient systems simulation, [39]). All variable thermodynamic environmental variables feeding the sensation models are determined within TRNSYS in steps set at one-hour increments. Consequently, the decision model can make no sub-hourly decisions. The sensation models and the instance-based decision model were coded in C++ and integrated into TRNSYS as so-called “types.”

The space that was modeled in TRNSYS for the parameter study consists of an office with a south-facing window, lacking a sunscreen and is located in Stuttgart. The external walls of the office were concrete with external insulation and the internal walls were modelled as plasterboard walls facing rooms with identical characteristics as the test room. The floor was a concrete slab with a floating screed and suspended ceiling. The space was occupied by one person, who was the only source of air pollution. At its current developmental state, no human occupancy model is integrated into the simulation. Consequently, the space is occupied 24 hours per day throughout the entire year, which is not in line with observable human behavior. However, this restriction was not deemed to conflict with the purposes of this study. Development and integration of an occupancy model are subject to future development.

The general model task aims to provide acceptable thermal environments and IAQ, meaning a sensation rating of zero on the respective sensation scale.

To achieve its thermal and air quality sensation goals, the model was provided with three action dimensions: the operation of the windows and heating systems and the adjustment of the clothing. Each of these dimensions are comprised of four distinct available action space states, listed in

Table 1.

Table 1: Available action dimensions and space states.

| Space states | Dimensions | | |
|----------------|---------------------|---------------------|--------------------|
| | Dim _{wind} | Dim _{heat} | Dim _{clo} |
| | [m ²] | [kJ/h] | [clo] |
| State 1 | 0.00 | 0.0 | 0.5 |

| | | | |
|----------------|------|----------|-----|
| State 2 | 0.05 | 2,500.0 | 0.8 |
| State 3 | 0.10 | 5,000.0 | 0.9 |
| State 4 | 0.40 | 10,000.0 | 1.1 |

The state of each of the action dimensions is part of the perceivable context for the model.

2.3 Procedure of the cognitive parameter study

2.3.1 Scope and sequence of studied parameters

The decision and learning model comprises a set of cognitive parameters. This study seeks to achieve a parametrization suitable for modelling energy-relevant behaviors in buildings. These parameters include:

- The similarity parameter SIM and the corresponding relaxation factor RF , which denote human abilities to discriminate the environmental cues relevant for action.
- The maximum associative strength S , which is used to determine the cognitive activation of memory instances through the perception of contextual cues before a decision.
- The noise parameter σ , which introduces an element of indeterminacy to the process of memory retrieval and is responsible for the balance between exploration of the world and exploitation of known solutions to particular problems.
- The partial matching parameter pm , which finetunes the context-specific activation of memory instances based on the degree of fittingness between the perceived context and the information stored in the memory instances.
- The decay parameter d , which rules the activation of memory instances through repeated execution of an action, or conversely, the degree of disremembering of unused information.

Each of these parameters and its meaning for the model will be introduced theoretically before testing. When available, empirical data will be used to substantiate the test spectrum for each parameter.

The parameters were tested in the sequence presented above. Simulations were carried out by varying only a single parameter at time, while setting all the remaining parameters to their default value or neglecting them. This is a heuristic approach that does not guarantee an optimal parameter setting, as the parameter settings might influence each other, which would not be covered by a linear testing procedure. Nonetheless, this approach was determined to be an adequate compromise between the need for accuracy and an appropriate effort.

2.3.2 Evaluation criteria

Parameter studies require the definition of a criterion against which the effect of parameter variations can be evaluated. Although individual empirical action-to-action data are an ideal criterion for validation, two reasons prevent this step: First, the lack of an occupancy model and thus the unrealistic ability of the model to act permanently, even at night or during periods of absence. Second, the current empirical data on occupant behaviour do not include all the relevant (psychological and sociological) details necessary to set all model parameters. Therefore, it was necessary for the validation process to develop criteria that are independent of concrete empirical data and are rather based on logical and common-sense considerations. These are presented in the following paragraphs.

Standard deviation of non-action prediction

An appropriate criterion for the optimization of the cognition-related parameters should ideally be objective, meaning that it should be independent of external data and serve as a reliable measure of performance across different contexts and situations. Thus, the *internal predictive consistency* of the model was the first criterion developed for testing the performance of the model. For a better understanding of the criterion, recall that the model has three action dimensions, each of which encompasses four potential action states in the test configuration (see Table 1). Thus, a total of 12

action states are available to the model during the decision process. For each decision, the model independently determines the effect of each of the 12 available states. Explicit non-action options are not included among the available action options; nonetheless, the model can choose to remain passive by not altering the current states of the elements. Thus, three of the 12 calculated effects (corresponding to three action dimensions) allow the model to remain passive, and these values independently predict future sensation states for identical circumstances. The central assumption of the *internal predictive consistency* criterion is that there is no reason for these three predictions to differ substantially if the model has the opportunity to gather sufficient experience with the entire spectrum of actions and conditions. As an example, if the window opening level is 0.05m², the heating emission is 5000 kJ/h and the clothing value is 0.9 Clo in a decision situation, then the three independent predictions for maintaining these states should ideally be identical because a human actor would likely not make a distinction between these three *non-action predictions*.

To measure the degree of similarity between the three non-action predictions, their standard deviation was calculated at each decision, and the frequency distribution of the standard deviation across one year for each type of sensation was used to compare different parameter runs. As in previous models, the time resolution was set to one hour such that the model makes one prediction per simulation hour. In total, each simulation ran three times over five years per parameter variation and only the results of the last year were considered for the study. The frequency distributions were averaged across the three simulation runs.

Apart from numerical imprecision, a model should ideally be able to produce highly consistent non-action predictions across the three dimensions, which would be in line with expectable human performance. This criterion is therefore considered to be both a *relative criterion* to compare the model performance with different parameter settings and an *absolute criterion* evaluating the absolute performance of the model.

Impairment vs. improvement of thermal conditions in winter after action

An additional criterion was introduced in case the first criterion does not yield unambiguous results. For this analysis, the thermal sensation *after* a decision was compared to the thermal sensation *before* a decision that was thermally elicited, and the results were separated into those decisions that led to an improvement and those that led to further impairment of the thermal conditions. The scope of this analysis was limited to the winter, during which appropriate means to counteract uncomfortable conditions are available to the model (specifically heating and clothing). Thus, this limitation restricts any effects that would result from a lack of appropriate action dimensions to counteract warm conditions during the summer, which would not be attributable to a lack of model performance.

Though it is a plausible assumption that decisions resulting in improved conditions are indicative of effective model performance, whereas decisions that result in impaired conditions are indicative of poor model performance, this criterion is not used as an absolute measure of the quality of the model. A decision that yields impaired conditions could be due to a misprediction (a bad model decision); however, it might rather be the best possible result of an appropriate decision that avoids even worse conditions. The analysis revealed that the model sometimes made decisions without the explicit goal of improving the environment, but rather to avoid further impairment of conditions.

Despite the restriction, this criterion was expected to be an appropriate supplementary indicator of the *relative* performance of the model.

Average and standard deviation of thermal conditions in winter

One of the goals of the model is to establish neutral thermal conditions. Though a lack of a cooling system makes it impossible for the model to achieve this goal during conditions of hot external temperatures, the provided action space theoretically enables it to be successful during winter. A further means to evaluate the *relative* performance of the model is, therefore, its ability in winter conditions to achieve, on average, neutral conditions.

3. Cognitive parameter study

3.1 Similarity and relaxation

3.1.1 Introductory theoretical considerations of the similarity parameter (SIM) and the relaxation factor (RF)

Human perception of the environment can be considered an essential pre-condition for planning and determining appropriate actions. As described in Section 2.1, the action model uses the PMV to predict thermal sensation and a comparable concept for predicting IAQ. This approach is a simplification; both types of models aim to predict sensation on the aggregated level of a larger group and are therefore not understood as representations of individual sensations. Even more importantly, these models predict sensation with an unrealistic accuracy of multiple decimal places. Simply importing this type of model into the action model would enable it to discriminate sensation states to a degree that is not in line with human perceptive abilities, such as the discernment of temperature fluctuations of 0.1 K, and the model would, in turn, fine-tune its actions with an unrealistic degree of detail.

The similarity parameter (SIM-value) accounts for realistic human perceptive abilities and is psychologically reducible to the concept of just noticeable difference (JND). First described by Weber [40], the JND describes the *objective* modality-specific difference between a reference stimulus and a comparison stimulus that can just be detected by a person. This concept includes the limitation that if two stimuli are too similar to each other, i.e., their difference is below the JND, they are usually no longer distinguishable, but rather are perceived as being identical. Consequently, such similar experiences must not be distinguished by the model during its interaction with the environment.

The SIM-value is relevant for three aspects of the model: 1) the integration of a new chunk into the declarative memory based on its similarity with already existing chunks; 2) the process of spreading activation from the encountered context through the declarative memory based on the similarity with the slots of the memory chunk, and 3) the perception of the environmental changes provoked by an executed action. During the first two of these processes, instance slots are compared to find identical slots. Treating sensation as a continuous rather than discrete space and comparing sensation slots to the smallest decimal place would lead to the accumulation of an infinite number of seemingly different instances in memory, which would be not only computationally expensive but also unrealistic from a cognitive perspective. Applying the JND concept through a SIM-value treats sensation as a discrete (rather than continuous) state, which results in a significantly more efficient and realistic process. Consequently, a SIM-value need not be defined for instance-slots that are discrete by nature, such as the clothing value or the state of the window. In the third process, the model senses environmental changes between the pre-and the post-action sensation states and thus establishes context-action-result contingencies. However, if the difference falls within the plus/minus range of SIM-value relative to the initial state, the environmental change is assumed to be undetectable and is therefore nullified.

The process of discretizing the space for thermal sensation initiates with the neutral sensation, such that zero is the center of the first category, which expands by $SIM/2$ in a negative or positive direction. Further categories are established with center points corresponding to multiples of the SIM -value, which ensures that deviations from zero of the same absolute magnitude are categorized symmetrically for positive and negative sensations. Asymmetrical scales such as the air quality sensation scale are discretized in a parallel manner, resulting in the first (neutral) category being half the size of the SIM -value. Fig 2 illustrates the principle.




Fig 2: Discrimination principle for sensation spaces.

The red “x” in Fig 2 illustrates the principle that each new experience produced by the sensation models with high numerical accuracy belongs to one of the pre-defined sensation categories and is then categorized accordingly. However, the two “xs” that are very close to each other on the positive side demonstrate an implausible consequence of this classification process, whereby a sharp and deterministic separation between two almost identical sensations located close to a category boundary always leads to these sensations being categorized differently, no matter how small their differences. This situation was assumed to be unrealistic, and therefore, the application of a random element for the categorization process was tested, thus reflecting the more realistic assumption that humans would not make such sharp distinctions between such similar sensations.

This random element is established by a sigmoid function applied to the boundary between two sensation categories. Equation (1) illustrates the format of the function for the example of the PMV-value produced by the thermal sensation model. This equation determines the probability with which a PMV-value is assigned either to the appropriate range of category CAT or to the adjacent category, and it includes a relaxation factor RF , which determines the steepness of the sigmoid curve and thus the sharpness of the distinction between the two involved categories.

$$prob = \frac{1}{1 + e^{\pm \left(PMV - CAT \mp \frac{SIM}{2} \right) \cdot RF}} \quad (1)$$

Fig 3 compares the probability distributions for relaxation factors 30 and 100 at a SIM -value of 0.4. For both factors, there is a 0.5 probability of a sensation at the boundary between two categories (in the example, -0.2 and 0.2) being assigned to either category. Close to the boundary, the gradient

determines the range in which the assignment is ambiguous, which is significantly smaller for a relaxation factor of 100 than for 30, thus indicating a more deterministic assignment.

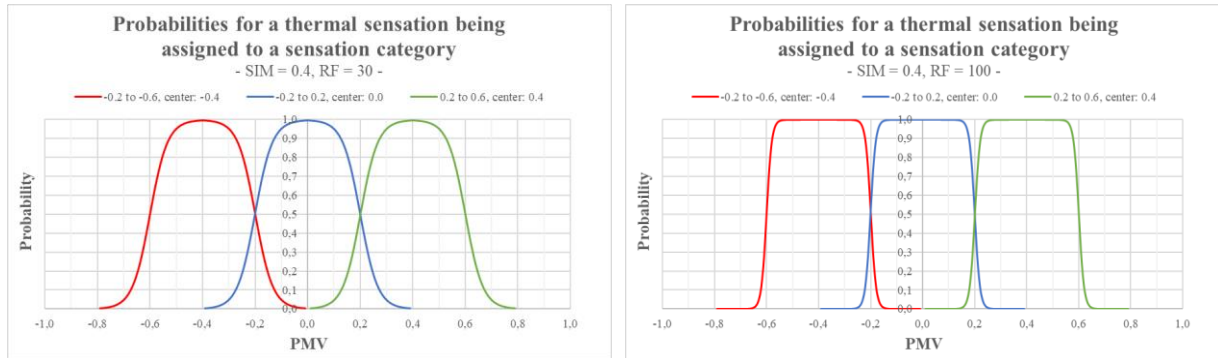


Fig 3: Probability of a continuous sensation to be assigned to a particular sensation category as dependent on the relaxation factor.

Previous versions of the model [28, 38] not only used a SIM-value, but also implemented a perception threshold “around” the goal state, thus rendering the model passive in case the sensation fell into the range defined by the perception threshold. Establishment of a perception threshold thereby also reflects the idea that slight deviations from the goal state are not perceptible and therefore will not elicit actions. This notion is based on the same considerations as those underlying the SIM-value, which is the just noticeable difference. However, previous models parametrized the SIM-value and the perception threshold independently of each other, which contradicts their shared theoretical origin. Therefore, as explained above, the current model did not treat these parameters separately but rather defined a single SIM-value.

The determination of the SIM-value for this model is based on the experiments of Collins et al. [41] and Natsume et al. [42]. In their experiments they found discriminable temperature differences depending on age (Collins/Natsume: 2.4/0.9°C (young males), 4.8/1.6°C (old males)) which differ significantly. However, if these values are converted into PMV-based thermal sensation values based on the information given in the studies, a consistent picture emerges with values 0.45 (Collins et al.) and 0.47 (Natsume et al.) among the young participants and 0.90/0.83 for old (healthy) subjects. For subsequent simulations, the similarity value for thermal sensation was set to 0.45 scale points in alignment with the data for young, healthy subjects.

These experiments provide information about the discriminative abilities of the model at conditions around the neutral thermal sensation. However, the characteristics of the JND indicate that these capabilities are constant across the entire range of sensations if measured on the sensation scale. Usually, the JND refers to objective, measurable stimulus magnitudes and increases proportionally to the magnitude of the reference stimulus S_m ($JND = \Delta S_m = k * S_m$, with k being the Weber constant). Although the objective just noticeable stimulus difference ΔS_m changes with stimulus intensity, the corresponding perceptual experience is by definition equal for all JNDs (independent of their objective magnitudes). Since the SIM-value refers to the subjective sensation of the stimulus, the SIM-values are also assumed to be constant across the range of the sensation spectrum. Given the lack of any other evidence, the SIM-value is also assumed to be constant across modalities, such as for air quality sensation in this action model. Fig 4 re-illustrates the principle with reference to the above-explained experiments.




Fig 4: Applied discretization of the continuous sensation scales based on [41] and [42].

3.1.2 Effect of similarity and relaxation on model performance

A set of similarity values ranging from 0.25 to 0.55 was implemented in the model and the according changes of the model performance were analyzed. Based on previous simulations [28, 38, 43], the maximum associative strength S was set to 30. To minimize the effect of random, the noise parameter σ was set to 0.05 and base level activation and partial matching was deactivated.

Fig. 5 illustrates a positive correlation between increasing SIM-values for thermal sensation and the ability of the model to make consistent predictions of the effect of non-action across all three action dimensions. This is indicated by an increasing proportion of low standard deviations (< 0.02) between the three non-action predictions (bins beyond 0.04 are not shown because they have no information value and to optimize readability). A one-way ANOVA revealed highly significant differences between the 0.02 bins of different SIM-settings ($F = 24.7/4.1$, $p < 0.01$); however, this effect was not observable for air quality sensation (i.e., the differences were not significant, $p = 0.14$).

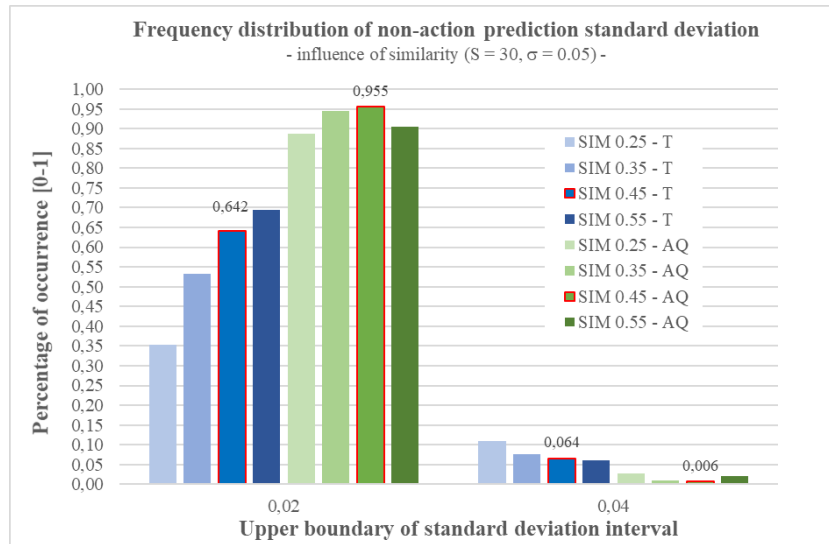


Fig. 5: Baseline for the distribution of standard deviations of non-action predictions in the thermal and air quality sensation dimensions.

However, what might seem to be improving the model performance in the thermal dimension is simply an effect of the degraded ability of the model to discriminate states. With rising SIM-values, more and more results are set to zero, which decreases the differences between the non-action

predictions. Thus, the frequency of the standard deviation of 0.02 and below of 0.642 for a SIM-value of 0.45 represents the baseline for the remaining parametrization process of the model.

Another obvious effect is that the performance was lower in the thermal dimension than in the air quality dimension, which was not unexpected because the window dimension is the sole action dimension that directly influences the sensation of the air quality; the model easily learns that clothing does not affect air quality and heating has very little effect (there is a small effect because changes of the internal temperature influence the buoyancy effect when the window is opened). Consequently, the uncertainty of the prediction relates only to the effect of the current window opening level. In contrast, all three action dimensions exert substantial influence on the thermal sensation, and the prediction in each dimension is thereby biased by some uncertainty. These uncertainties result in a greater deviation between predictions than in the case of air quality sensation.

The decision quality (impairment vs. improvement) was not significantly affected by SIM-values and is therefore not presented here.

Both the above-described experiments by Collins et al. [41] and by Natsume et al. [42] found that the groups of subjects established comparable *average* temperatures in the climate chamber irrespective of their ability to discriminate the temperature stimuli and the corresponding magnitude of the observable temperature amplitudes. Therefore, the average sensation and the standard deviation of the sensation established by the model for winter conditions were investigated for the analyzed SIM-values. Fig 6 presents the results for the established thermal sensation, and it can be observed that there is only a slight dependence of the average thermal sensation on SIM-value, which is in line with the experimental data. Additionally, as represented by the standard deviation SD, the wider expanse of the established conditions corresponds well with increasing SIM-values, which is also in line with expectations.




Fig 6: Comparison of established average thermal conditions and the corresponding standard deviations for all investigated SIM-values.

Tests were also conducted using a SIM-value of 0.45 to measure the effect of a set of different relaxation factors *RF* in the range of RF 30 to RF 300 on the model performance. Fig 7 illustrates the effect of applying a *RF* on the consistency of non-action predictions (data-labeling refers to RF = 140). As indicated by the larger fraction in the 0.02- category compared to the simulation without *RF* (red columns), there was a slight tendency of the performance to improve with increasing *RF*-values;

however, the improvement was inconsistent and statistically not significant ($F = 0.85/3.11$, $p = 0.54$ for thermal sensation). A one-sided t-test also revealed insignificant differences between the results achieved without relaxation vs. those yielded with an $RF = 140$ and an $RF = 300$ ($p_{140} = 0.27$, $p_{300} = 0.12$ for thermal sensation, respectively).




Fig 7: Comparison of the distribution of standard deviations of non-action predictions in the thermal and air quality sensation dimensions for different relaxation factors RF .

The analysis of the decision quality in winter conditions reveals a slightly clearer picture. An RF of 140 minimized the occurrence of situational impairment after a decision and maximized improvement accordingly, and the improvement was significant ($p = 0.03$) compared with the performance with no relaxation. In contrast, there was only a slight, insignificant improvement when using an RF of 300 ($p = 0.29$). Based on these results, an RF of 140 was used for subsequent steps in the parametrization process.




Fig 8: Effect of different relaxation factors on post-action impairment vs. improvement of thermal conditions in winter conditions.

3.2 Spreading activation

3.2.1 Introductory theoretical considerations of the spreading activation parameter S

In ACT-R, the spreading activation mechanism rules the context-driven retrieval of information from the declarative memory [30]. This concept assumes that chunks in declarative memory do not exist in isolation, but rather form an interconnected network. In a retrieval attempt, the contextual information at hand spreads activation through the network and activates chunks that match the context. The higher the activation of a chunk in memory, the more likely its retrieval [44].

The spreading activation SA received by a memory instance i is determined by the number of slots s in the instance representing the context and their fittingness with the memory instance slots, and it is calculated according to equation (2):

$$SA_i = \frac{1}{s} \cdot \sum_{j=1}^s S_{ji} \quad (2)$$

In this equation, j counts through all s slots of the source instance.

According to equation (2), the total spreading activation SA received by a memory instance i equals the arithmetic mean of the associative strength S_{ji} across all source slots s . The associative strength S_{ji} was initially defined as the ratio between the likelihood of retrieval $P(i/j)$ of chunk i in a cueing context j and the unconditional likelihood $P(i)$ [45]. Since then, the standard approach to determining the associative strength S_{ji} between source slot j and chunk i in memory has evolved, such that the log-likelihood ratio of $P(i/j)$ and $P(i)$ is used. $P(i/j)$ is currently approximated by $= 1/\text{fan}_j$, with fan_j representing the number of times the source slot j is part of any of the memory chunks (“fan”, [46, 47]). Thus, “fan” models the ambiguity of the contextual information for the retrieval of a chunk. $P(i)$ can be approximated with $1/n$, with n representing the total number of chunks in memory; however, its logarithm is usually set to a constant model parameter S , such that S_{ji} results in equation (3):

$$S_{ji} = \ln\left(\frac{\frac{1}{\text{fan}_j}}{\frac{1}{n}}\right) = S - \ln(\text{fan}_j) \quad (3)$$

In most applications, the maximum associative strength S is set to a value around 2.0 in ACT-R [30]. Since a fan higher than seven would already render the associative strength inhibitory, Thomson and Lebiere [48] pointed out that this might only be appropriate for modelling a single session psychological experiment with a limited number of relevant chunks in memory. However, fan is usually high in the present application of spreading activation, which entailed five-year simulations of energy-relevant interaction with buildings and for which some thousand chunks were created in a limited context. This would result in the inhibition of memory chunks for an S of around 2.0. Consequently, S was estimated to be as high as 30 for previous versions of this model ([28, 38, 43]). Although in combination with other settings, this value has been shown to lead to plausible model behavior, for the present model, the value was modified based on equation (4), in which S is determined based on the estimated number of chunks n in memory.

$$S_{ji} = \ln(n) - \ln(\text{fan}_j) \quad (4)$$

3.2.2 Effect of the spreading activation parameter on model performance

The previous simulations showed that the total number of chunks n in the declarative memory converged at around 2000 for actions related to thermal sensation and at approximately 1000 for actions related to air quality in the fifth year of the simulation. For a total of 4-5 context slots, this

yields 8000/4000 to 10,000/5000 (T/AQ) pieces of information stored in the memory, which results in an S ranging between $\ln(2000) = 7.6$ (if n is considered to be the number of chunks) and $\ln(10,000) = 9.2$ (if n is considered to be the number of slots) for thermal instances and $\ln(1000) = 6.9$ or $\ln(5000) = 8.5$ for air quality instances. Based on these considerations and to allow for some variance in the maximal number of chunks, an adequate maximum associative strength was estimated to be around a value of $S = 12$. For the sake of simplicity, S is considered to be a general parameter, rather than a modality-specific parameter, and thus is applied to thermal instances as well as air quality instances.

Consequently, the model performance was tested for a set of values for the maximum associative strength S ranging from $S = 7$ to $S = 25$. Fig 9 presents the results for the non-action prediction consistency. For values of $S = 25$ and lower, the performance seems to have been slightly weaker in comparison with a maximum associative strength of $S = 30$. However, as expected, an ANOVA revealed no statistically significant differences between different S levels ranging from 12 to 30 for either type of sensation, nor did one-sided t-tests reveal any significant differences when $S = 30$ was paired with any of the other strengths between $S = 12$ and $S = 25$.

In contrast, a maximum associative strength of $S = 7$ led to a substantial and significant (t-test, $p < 0.01$) performance impairment for thermal sensation in comparison with $S = 30$. This was also to be expected because the average $\ln(\text{fan})$ for each of the slots is usually in the range of 6-7 in such cases, thus rendering the total spreading activation close to zero, which almost neutralizes the effect of context on activation. Consequently, with this parametrization, the model tends to choose its actions independent of the prevailing context.

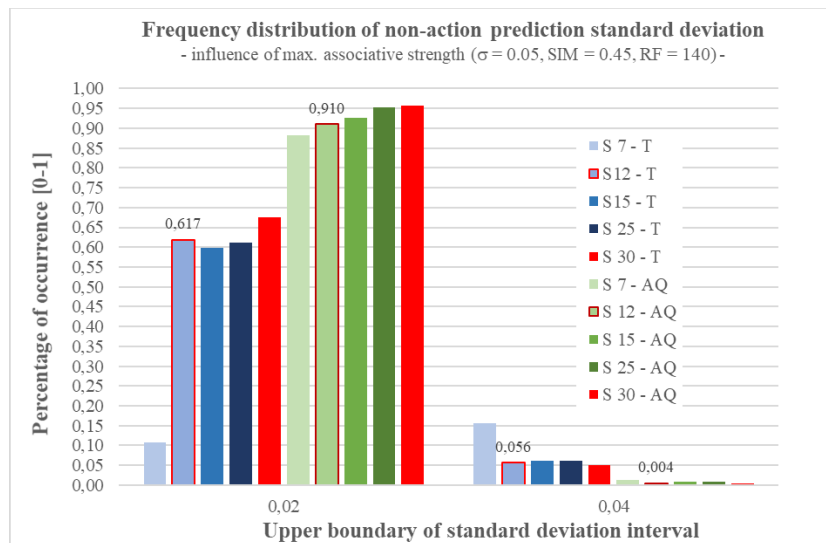


Fig 9: Comparison of the distribution of standard deviations of non-action predictions in the thermal and air quality sensation dimensions for different maximum associative strengths S .

Fig 10 illustrates the results of an investigation of the decision quality in winter conditions based on thermal sensation for S -values ranging from 12 to 30, which compared the relative advantages of a maximum associative strength of $S = 12$ with those of higher values. Whereas an ANOVA revealed significant differences between the different parametrizations ($F = 13.29/4.07$, $p < 0.01$), a one-sided t-test indicated no significant advantage for $S = 12$ over $S = 30$.




Fig 10: Effect of different maximum associative strengths on post-action impairment vs. improvement of thermal conditions in winter conditions.

3.3 Noise and its influence on model performance

3.3.1 Introductory theoretical considerations of the noise parameter σ

Indeterminacy in the instance retrieval process enables the model to occasionally make seemingly non-optimal decisions and thereby further explore the action-result-contingency space provided by the environment and broaden its knowledge of the world. A total absence of noise and indeterminacy would lead the model to resort to and exploit the “so-far best” solution, even if more appropriate solutions might be available. Conversely, excessive noise in the retrieval process would lead to random, arbitrary decisions and erratic model behavior. To achieve a more balanced behavior requires an optimal noise factor in the sense of the quality criterion introduced in before.

The arithmetic used to determine the influence of the noise parameter were entirely adopted from IBLT [49-51]. The influence of chance is governed by the noise parameter σ , which in turn influences two different mechanisms, namely the noise activation of memory instances N and the temperature factor τ in the softmax equation that governs the probability of retrieval P . At every retrieval request, every chunk in memory receives a noise activation that is governed by equation (5):

$$N_i = \sigma \cdot \ln \left(\frac{1-\gamma_i}{\gamma_i} \right) \quad (5)$$

In this equation, σ scales the effect of γ , which is a random draw from a uniform distribution between zero and one. Fig 11 illustrates the influence of σ and γ on noise activation.

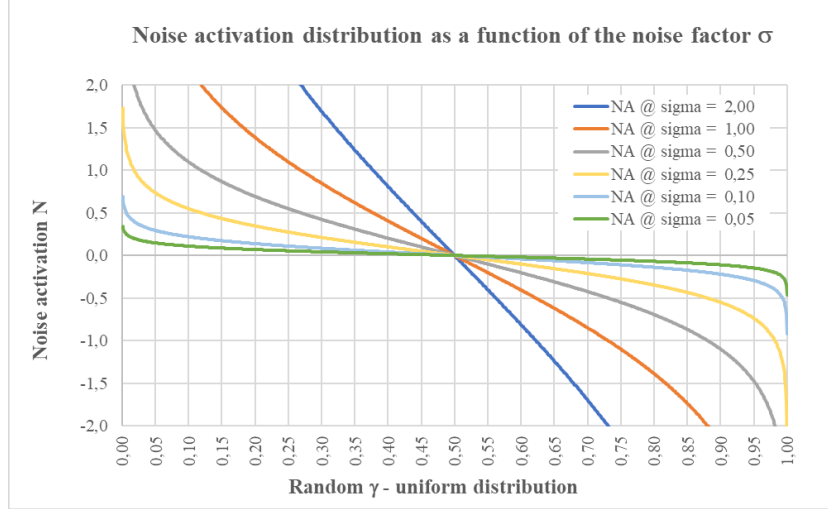


Fig 11: Influence of γ and σ on instance activation by noise.

Some analysis is required to understand the effects of σ and γ . In the retrieval process of an action, the retrieval probability P for a memory instance i is used as a weighting factor of the action effect stored in the instance, as represented by the utility U of this effect in the current situation. To determine the blended value BV of a particular action m – i.e. its expected effect – the product of retrieval probability and utility is summed over all instances r_m belonging to that action.

$$BV_m = \sum_{i=1}^{r_m} P_{i,m} \cdot U_{i,m} \quad (6)$$

The probability of retrieval of an instance i is determined by the softmax equation (7), in which τ is the temperature factor and which is determined by equation (8). For each instance i belonging to a particular action, the softmax equation determines the ratio between the exponential of the instance's activation A_i and the sum of all instances related to this action over the exponentials of the activation. This yields a probability distribution with unity sum.

$$P_i = \frac{e^{\frac{A_i}{\tau}}}{\sum_{j=1}^{r_m} e^{\frac{A_j}{\tau}}} \quad (7)$$

$$\tau = \sigma \cdot \sqrt{2} \quad (8)$$

The numerator of the softmax equation was further analyzed to estimate the effect of noise on retrieval probability. Specifically, the probability that noise activation would lead to an equal or higher activation of an instance with low spreading activation (i.e., weaker fittingness with the current environmental cues) was examined in comparison with an instance with high spreading activation (i.e., with a better fittingness with the current environmental cues). With spreading and noise activation as only sources of activation, the numerator of equation (7) can be rewritten according to equation (9):

$$e^{\frac{A_i}{\tau}} = e^{\left(\frac{SA_i}{\tau} + \frac{N_i}{\tau}\right)} = e^{\left(\frac{SA_i}{\tau}\right)} \cdot e^{\left(\frac{N_i}{\tau}\right)} \quad (9)$$

Combined with equations (5) and (8), this yields

$$e^{\frac{A_i}{\tau}} = e^{\left(\frac{SA_i}{\sigma \cdot \sqrt{2}}\right)} \cdot e^{\left(\frac{\sigma \cdot \ln\left(\frac{1-Y_i}{Y_i}\right)}{\sigma \cdot \sqrt{2}}\right)} \quad (10)$$

In equation (10), σ cancels out in the second factor such that the part of the retrieval probability related to noise activation only depends on γ and renders the according part of equation (10) a factor independent of σ . The noise factor σ only affects the influence of spreading activation by scaling the related “temperature.”

Two instances with different SA and γ were now compared to determine the probability at which the difference in spreading activation between two instances would be offset by their noise. In this calculation, SA_1 was higher than SA_2 and the necessary γ_2 was determined to give the second instance an equal or higher retrieval probability than the first.

$$e^{\left(\frac{SA_1}{\sigma \cdot \sqrt{2}}\right)} \cdot e^{\left(\frac{\ln\left(\frac{1-\gamma_1}{\gamma_1}\right)}{\sqrt{2}}\right)} \leq e^{\left(\frac{SA_2}{\sigma \cdot \sqrt{2}}\right)} \cdot e^{\left(\frac{\ln\left(\frac{1-\gamma_2}{\gamma_2}\right)}{\sqrt{2}}\right)}$$

Reordering yields:

$$\gamma_2 \leq \frac{1}{e^{(SA_1-SA_2) \cdot \frac{1}{\sigma \cdot \sqrt{2}} \cdot \frac{1-\gamma_1}{\gamma_1} + 1}} \quad (11)$$

Thus, the second instance receives a higher total activation than the first instance for all γ_2 lower than the right side of equation (11). Since γ follows a uniform distribution between 0 and 1, the value of the right side yields a probability of γ_2 , thus resulting in a retrieval probability of the second instance that is equal or higher than that of the first instance with its given γ_1 , a different spreading activation ΔSA , and a noise factor σ . Finally, with a given ΔSA and σ , integration over γ_1 between 0 and 1 yields a total probability $P(A)$ of the activation of the second instance that is equal or higher than that of the first instance.

$$P(A) = \int \frac{1}{e^{\Delta SA \cdot \frac{1}{\sigma \cdot \sqrt{2}} \cdot \left(\frac{1}{\gamma_1} - 1\right) + 1}} d\gamma_1 \quad (12)$$

Integration yields:

$$P(A) = - \frac{e^{\Delta SA \cdot \frac{1}{\sigma \cdot \sqrt{2}} \cdot \ln\left(\left|\left(1 - e^{\Delta SA \cdot \frac{1}{\sigma \cdot \sqrt{2}}}\right) \cdot \gamma_1 + e^{\Delta SA \cdot \frac{1}{\sigma \cdot \sqrt{2}}}\right|\right)} + \left(e^{\Delta SA \cdot \frac{1}{\sigma \cdot \sqrt{2}}} - 1\right) \cdot \gamma_1}{\left(e^{\Delta SA \cdot \frac{1}{\sigma \cdot \sqrt{2}}} - 1\right)^2} \quad (13)$$

Fig 12 illustrates the results of equation (13) for a set of σ and ΔSA values. It can easily be observed that a noise factor of 0.05 leads to a very strict separation of instances with an almost zero percent probability of instances with a difference in spreading activation ΔSA of 0.5 having the same retrieval probability. In contrast, a noise factor of 0.50 leads to a substantial relativization of the influence of context on retrieval probability because the retrieval probability is the same or even reversed for almost 35 % of those instances with a ΔSA of 0.5.

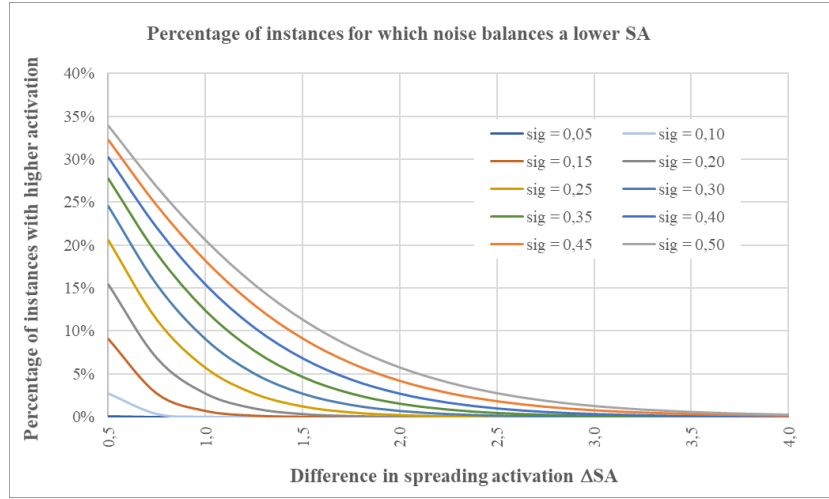


Fig 12: Percentage of instances that would receive a higher total activation through the effect of noise than other instances with higher spreading activation.

To put these values into perspective, the average activation of instances was estimated based on the simulation results previously produced for this study. These simulations showed that a fan of around 400 could be expected averaged across all decisions during the last simulation year and across all four context slots, which is roughly equivalent to a penalization of the maximum spreading activation by a value of 6.0. With a maximum spreading activation of $S = 12$ (as determined in section 3.2), a net activation of $12 - 6.0 = 6.0$ was achieved in cases when all four context slots matched the source slots. Thus, roughly speaking, every slot is responsible for activation of $6.0/4 = 1.5$. Based on this estimate, the relativization of contextual influence as caused by a noise factor of $\sigma = 0.5$ seemed intuitively to be too high; rather, a factor between 0.1 and 0.25 was expected to be an adequate choice.

3.3.2 Effect of noise on model performance

With base-level activation and partial matching still deactivated, the model performance was tested to compare the impact of different σ values. The results are summarized in Fig 13, which shows a considerable performance improvement for thermally driven decisions in comparison with Fig 9 and a σ value of 0.05 for σ values in the range of 0.07 to 0.12. Beyond these σ values, however, the performance drastically degrades below the values shown in Fig 9. However, despite these large differences, t-tests comparing the performance with $\sigma = 0.05$ to those with σ values of 0.07, 0.10 and 0.12 revealed that these differences were not significant ($p > 0.20$ in each case). In line with this, the performance difference between σ values of 0.12 and 0.15 was also not significant ($p = 0.17$). However, the differences between σ values of 0.12 and σ values of 0.20 ($p = 0.03$) and 0.25 ($p < 0.01$) were significant.

The effect of σ was considerably less for the olfaction related predictions, although a comparable tendency is recognizable. No statistically significant differences exist between the performance for σ values of 0.12 and 0.20 (one-way ANOVA, $F = 1.60/3.15$, $p = 0.23$).

Thus, the testing of different σ -values did not reveal an unambiguous picture for the range close to $\sigma = 0.05$, as it remains unclear whether σ -values between 0.07 and 0.15 improve the performance. However, performance in the testing conditions is impaired for $\sigma \geq 0.20$, and testing of the decision quality in winter conditions (impairment vs. improvement) revealed no further insights. Since the

theoretical explanations in section 3.3.1 and Fig 12 indicated the need to choose a higher σ -value, further simulations used the highest σ -value with the best performance, which was $\sigma = 0.12$.

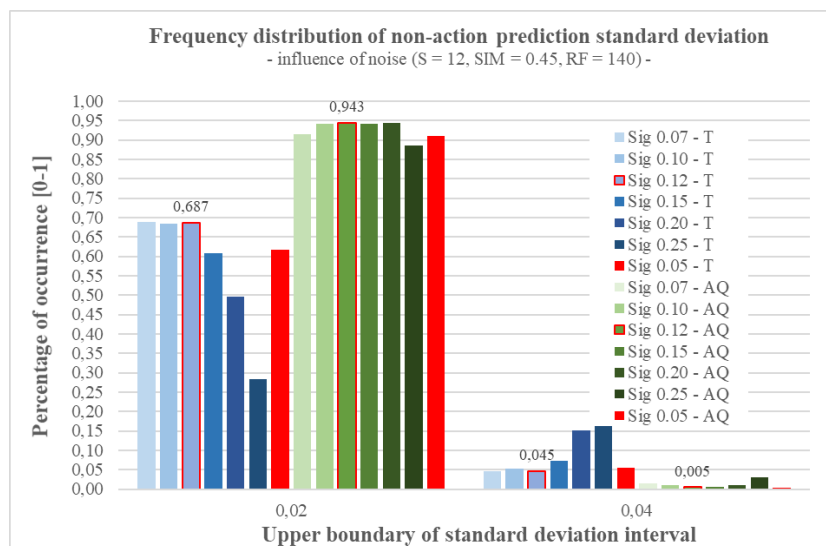


Fig 13: Influence of noise on model performance.

3.4 Partial matching

3.4.1 Introductory theoretical considerations of partial matching

Consider the question “How many animals of each kind did Moses take on the ark?” In an experiment, nearly half of the participants answered with “two” without realizing that it was not Moses’ but Noah’s ark [52]; this effect is called the “Moses illusion.” In contrast, people are not usually tricked by the question “How many animals of each kind did Nixon take on the ark?” [53]. The relevant difference is the semantic relatedness between Moses and Noah and the lack of the same between Nixon and Noah. Thus, while trying to retrieve the correct answer from memory, the process tends to tolerate slight mismatches in the question so long as they are close enough to each other. This is a common effect in human cognition and is co-responsible for the ability to generalize across conceptually comparable contexts. As such, memory retrieval enables sensible and goal-directed behavior in partially unfamiliar conditions; however, this process might also be the source of errors.

In previous simulations, only those instance slots that exactly matched the context contributed to the spreading activation of a memory instance, whereas non-matching slots were ignored and instances with no matching slots received a zero activation. However, in modelling energy-relevant interaction, it also makes sense to consider spreading activation not as a dichotomous variable, but rather as a continuous variable depending on the degree of slot fittingness. For example, at a clothing state of 0.8 clo, it is relevant to distinguish between a memory instance slot value of 0.9 and one of 0.5. In the first case, the context almost matches and the prediction for maintaining the clothing state (or actively deciding for a clothing state of 0.8 Clo) can be expected to be close to reality; although such an instance does not exactly match the clothing state slot, it would still be a valuable source of information. However, the context is significantly different in the second case, and the stored result (i.e., an anticipated increase of thermal sensation after changing the clothing from 0.5 clo to 0.8 clo) would not coincide with the true effect. However, the model could not make distinctions between these memory slots without considering the degree of fittingness.

In ACT-R, the concept of partial matching refers to the degree of fittingness between the slot information sought and the slot information that is found in the declarative memory. For example, to model the “Sugar factory” problem, a task in which participants must control a dynamic system with unknown and partly randomized rules [54], Taatgen and Wallach developed an instance-based ACT-R model that used partial matching [55]. This feature enabled the model to find valid solutions in line with experimental data by retrieving solutions from memory that were sufficiently close, albeit not an exact match to the currently experienced situation. The ability to generalize is beneficial in this context; however, generalizing can also lead to the retrieval of wrong solutions, specifically in cases when solutions must be exact, as in numerical operations. For example, Siegler and Shrager investigated the addition strategies and abilities of 4- and 5- year-old children for low numerical values, such as $1 + 5$ [56]. Strikingly, the children’s wrong answers were distributed around the correct answers, and there was a tendency for the retrieval probability of a wrong answer to decrease with a greater numerical difference to the correct answer. Thus, the probability of retrieving the wrong result $1 + 5 = 5$ or $1 + 5 = 7$ was higher than the retrieval probability for the result $1 + 5 = 9$. These results align with the concept of scaling the activation of memory instances by the degree of fittingness between the contextual information at hand and the information stored in declarative memory, and they can be modelled in ACT-R by partial matching.

Partial matching PM in the presented model was governed by equation (14):

$$PM = -pm \cdot \text{abs}(\Delta\text{slot_value}) \quad (14)$$

Thus, the implementation of partial matching will not change the activation resulting from exactly matching slots. However, the spreading activation for non-matching slots will now not be set to zero; rather, these slots will receive an additional *negative* activation, potentially resulting in an inhibitory effect for these instances. The penalization is scaled by the pm -factor and is proportional to the absolute numerical difference between the slot values that represent the currently experienced situation and the slot values stored in the memory instance.

To estimate the magnitude of the pm -factor, its effect was studied for a concrete example before running multiple simulations. For this example, the declarative memory produced by one of the previous simulations with a noise factor of $\sigma = 0.12$ was extracted and used as the basis for the prediction of the effect of deciding for a Clo-value of 0.5, 0.8 and 0.9 Clo in an arbitrarily chosen situation defined by the action states for window (0.1 m^2), heating (5000 kJ/h) and clothing (0.8 Clo) and the thermal sensation states -0.90 , -0.45 , 0.00 , 0.45 and 0.90 . The clothing action was chosen because its effect on thermal sensation can be precisely predicted by the PMV-algorithm, which builds the basis for the perceived “environmental response” to this action. Although the model is not expected to make an exact prediction of the thermal sensation change provoked by adjusting the clothing, the exact values are used as a benchmark against which the quality of the predictions of the model is compared.

None of the context slots is supposed to receive more attention than the others. Thus, the pm -value of each context slot was chosen in such a manner that the activation penalization of each slot was roughly of the same magnitude if multiplied with the maximum action state difference. Since thermal sensation states did not span the entire spectrum of the sensation scale (-3 to 3) in the previous simulations, the pm -value for sensation was estimated based on a sensation span of -1.5 to 1.5 . The pm -values are summarized in Table 2.

Table 2: Partial matching pm-values for context slots.

| | pm-value | max. state difference | max. penalization |
|-----------|----------|-----------------------|-------------------|
| Window | 2.5 | 0.4m ² | 1 |
| Heating | 0.0001 | 10,000kJ/h | 1 |
| Clothing | 1.65 | 0.6 clo | 0.99 |
| Sensation | 0.3 | 3.0 | 0.9 |

The effect of the pm-value on predictive quality was then tested by multiplying all pm-values by factors ranging from zero (no partial matching) to 10.

Fig 14 shows the predictions for the three clothing adjustments across thermal sensation states ranging from -0.90 to 0.90. The red graphs are the exact values as calculated based on the PMV algorithm. Except for the action 0.8 Clo—which is a decision not to act, the effect of adjusting the clothing depends on the current thermal sensation. Furthermore, decreasing the clothing level from 0.8 to 0.5 leads to a cooler sensation, whereas an increase in clothing level from 0.8 to 0.9 produces a slightly warmer sensation.

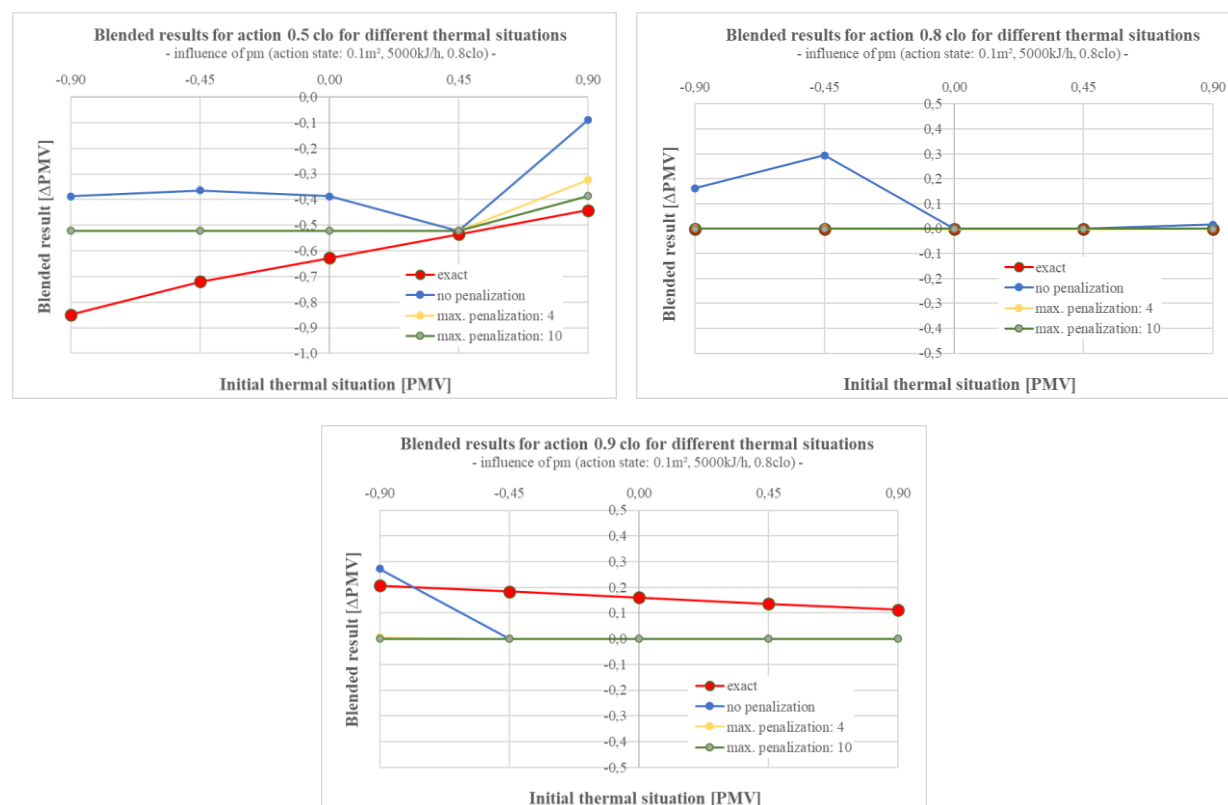


Fig 14: Model's predictive quality for changing Clo value in different thermal conditions and for different maximal penalizations.

The blue graph shows the predictions of the model without partial matching, and substantial deviations from the exact results can be observed specifically for 0.5 Clo. For the other actions, the deviations are less, albeit still noticeable. For reasons of clarity, only the results for a maximum penalization of 4 and 10 are shown in the figures. A clear improvement of the prediction quality can be observed for an action of Clo = 0.5; however, a maximum penalization of 10 yields no substantially improved result in comparison to a maximum penalization of 4. The predictions are not exact, which

is in line with the expectations and is not considered to be a model deficit, but rather representative of realistic, human-like performance.

For an action of 0.8 Clo, partial matching produces results that coincide with the exact values, meaning that the thermal sensation conditions do not change by the action, which is also in line with plausible human abilities; however, the model erroneously also expects an action of 0.9 Clo to not affect thermal sensation. Again, this is not considered to be a model deficit; rather, two previously discussed effects can presumably be responsible for such a result. First, the exact changes of sensation calculated based on the PMV-algorithm are below the similarity threshold of 0.45 points on the sensation scale and hardly detectable for the model. Such negligible changes are below the JND and are therefore stored as zero changes to the memory. Second, due to the close numerical similarity between Clo values of 0.8 and 0.9, memory instances with clothing slots of 0.9 are not substantially penalized because they are considered to be sufficiently similar to the experienced context. Such situations thereby align with the Moses illusion and are not considered to be representative of a model deficit.

3.4.2 Effect of partial matching on model performance

Based on the results elucidated in the previous section, the effect of a maximum penalization ranging from 2 to 12 on the model performance was tested with deactivated base-level activation.

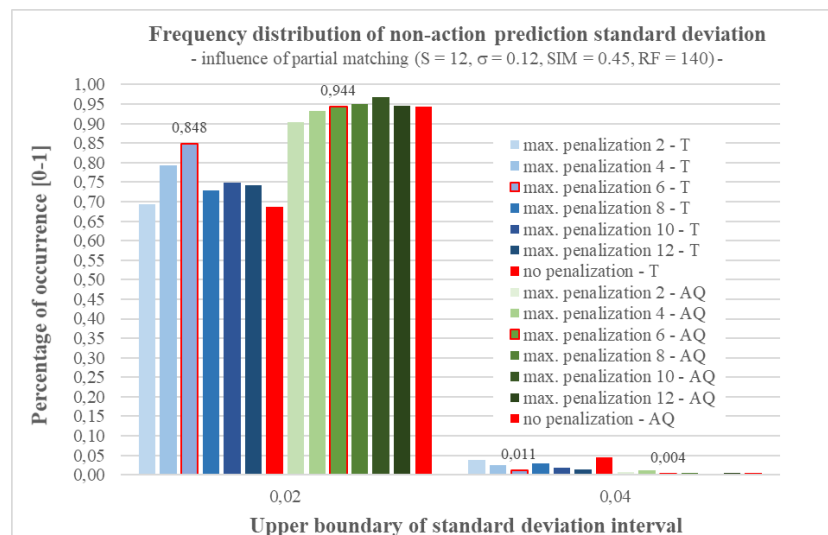


Fig 15: Influence of partial matching on the model performance.

For both criteria, namely non-action prediction consistency and decision quality in winter, partial matching with a maximum penalization of six produces the best results. In the case of thermal sensation, the percentage of non-action predictions with a standard deviation below 0.02 increases from 0.687 without partial matching to 0.848 with partial matching. Based on a one-sided t-test, the difference between the results without partial matching and those with a maximum penalization of six is significant ($p = 0.04$), whereas none of the other results with partial matching differs significantly from the version without partial matching. Again, the influence on the quality of the non-action prediction for air quality sensation is negligible and differences are not significant, presumably because the quality is already high without partial matching. Additionally, the quality of the decisions with partial matching based on thermal dissatisfaction in winter conditions is highest for a maximum

penalization of six and slightly better than without partial matching. The difference for a maximum penalization of six in comparison with the version without penalization is almost significant ($p = 0.08$).




Fig 16: Influence of partial matching on decision quality in winter conditions.

3.5 Decay and forgetting

3.5.1 Introductory theoretical considerations of decay and forgetting

Memories decay with time if they are not re-actualized. This effect is usually modelled in ACT-R by determining the base level activation B of chunk i according to equation (15).

$$B_i = \ln\left(\sum_{p=1}^e t_p^{-d}\right) \quad (15)$$

In this equation, the influence on the activation of a single encounter p decays with the time t_p that has passed since its occurrence, whereas the total activation is the sum over all past encounters e . The speed of decay is modified by the decay factor d , which is most often chosen to be 0.5 in ACT-R models [30]. Activation represents the accessibility of memories, and the underlying memory activation mechanisms aim at making most salient those memories that are most likely to be useful in a particular situation [57]. Whereas spreading activation determines the availability of a memory based on the currently given environmental cues, base-level activation refers to the effect of the recency and frequency of past occurrences of memories. Anderson and Schooler were able to demonstrate that this memory mechanism is well adapted to environmental characteristics and in fact, represents a correct statistical inference of environmental events [57]. Base level activation can therefore be considered a mechanism to improve the efficiency of memory processes since it does not rely on the perception of the current context.

Base level activation is an efficient memory mechanism in a patterned environment, but it also dampens or masks the influence of a more or less abruptly changing context. Depending on the speed of memory decay, activation displays inertia, which leads to a behavior lag, as demonstrated by Rakow and Miler for dynamic binary choice problems [58], among others. In Rakow and Miler's [58] experiments, participants' change of preference lagged significantly behind the change of objective payoff, thus leading to a non-optimal total reward.

A relevant but only rarely discussed aspect of equation (15) is the unit of the time t . The unit of time t and the magnitude of the decay factor d must correspond, meaning that d cannot be determined

without the determination of the time unit (or the other way around). Often, time is measured in trials or task blocks [49, 59]; however, this strategy is only applicable to experimental settings rather than real-world tasks. Others have used clock time between trials and “psychological time” (1/40 of real-time) between sessions as measures of time [60, 61]. Considering the irregularity of building occupant interaction patterns, an event-based measurement of time is certainly inappropriate because the time that passes between interaction events might differ by several hours in human-building interaction. Therefore, a clock-time-based approach is likely to be more appropriate. Additionally, using psychological time might be an adequate method to modelling the effect of decay during absence from the modelled space; however, testing this assumption requires modelling of occupancy times and times of absence, which is beyond the current capabilities of the model. Previous versions of the model used one hour as the time unit in combination with a decay factor of 0.5 [38, 43]. However, determining which decay factor is the most appropriate will be analyzed in more detail in this study.

Given the distinctive actions and types of action-result contingencies throughout a year in energy-relevant human-building interaction, memory decay likely plays a relevant role in human behavior in buildings. For example, certain types of equipment elements generally remain unused for weeks or months at a time, such as a heating system in summer or a cooling system in winter. Furthermore, the effect of an action such as adjusting the window depends on the seasons because the character of the weather and its influence on the internal conditions vary during different seasons of the year.

However, the empirical data indicating that a behavioral lag occurs in human-building interaction is weak and often only anecdotal. Additionally, unlike binary choice problems, in which events are independent of each other and success depends on both the objective payoff probability and a random process that determines the payoff in a concrete situation, success in energy-relevant occupant behavior is not independent of previous successes. Even if an actor operates an element based on the experienced successes of the past, the feedback to this concrete action is reliable, which leaves less room for interpretation than the binary choice problem. Wrong decisions can quickly be corrected with certainty, such as if opening the windows leads to unexpected and surprisingly cool indoor temperatures. This effect might mask the influence of base-level activation on real occupant behavior and thereby might be less easily detectable in human interaction data, specifically if only statistical macro-data instead of action-to-action data is available. If this assumption is true, a behavioral lag should be more obvious in situations in which constraints hinder immediate correction of wrong decisions, as Morgan and de Dear argued often occurs for clothing behavior [62]. They investigated the dependence of clothing behavior on indoor and outdoor temperatures in an office environment and a shopping mall [62]. In both cases, the decision for a clothing ensemble was made before arriving at the space in question, which restricted the available action space. In line with this argumentation, Morgan and de Dear [62] demonstrated a significant statistical relationship between the clothing level in the shopping mall on day x and the average daily outdoor temperatures of day x and the seven preceding days. This general pattern was confirmed by the study of de Carvalho et al. (among others [63]).

These data indicate that base level activation impacts behavior for around 1-7 days, which is equivalent to 24-168 hours. In combination with the typical interaction frequency of the model in each dimension, the magnitude of the decay factor can thus be narrowed down to a range that can then be tested through simulation. Over one year of the interaction of the model, the previous simulations revealed an average frequency in the range of seven interactions per 24 hours, which were evenly

distributed across all three action dimensions. Therefore, 2.0 to 2.5 interactions occur per day per dimension, thus resulting in several interactions per action state ranging from around 1.25 (2.5/2, switching between one particular state and any of the other states) to 0.5 (2.0/4, switching through all states).

Fig 17 presents an estimate of the base level activation for decay rates of $d = 0.30$ and 0.60 as a function of the number of actions per day (and the according activation) and the amount of time passed since the action was executed the last time. The comparison indicates that for most of the cases, a decay factor of $d = 0.30$ enables actions to remain activated even after they have not been executed for a week. In contrast, a decay rate of 0.60 leads to a rapid inhibition of actions within less than a day. Consequently, it was determined that a decay rate between these two values should be appropriate to model the observed lag in clothing behavior.

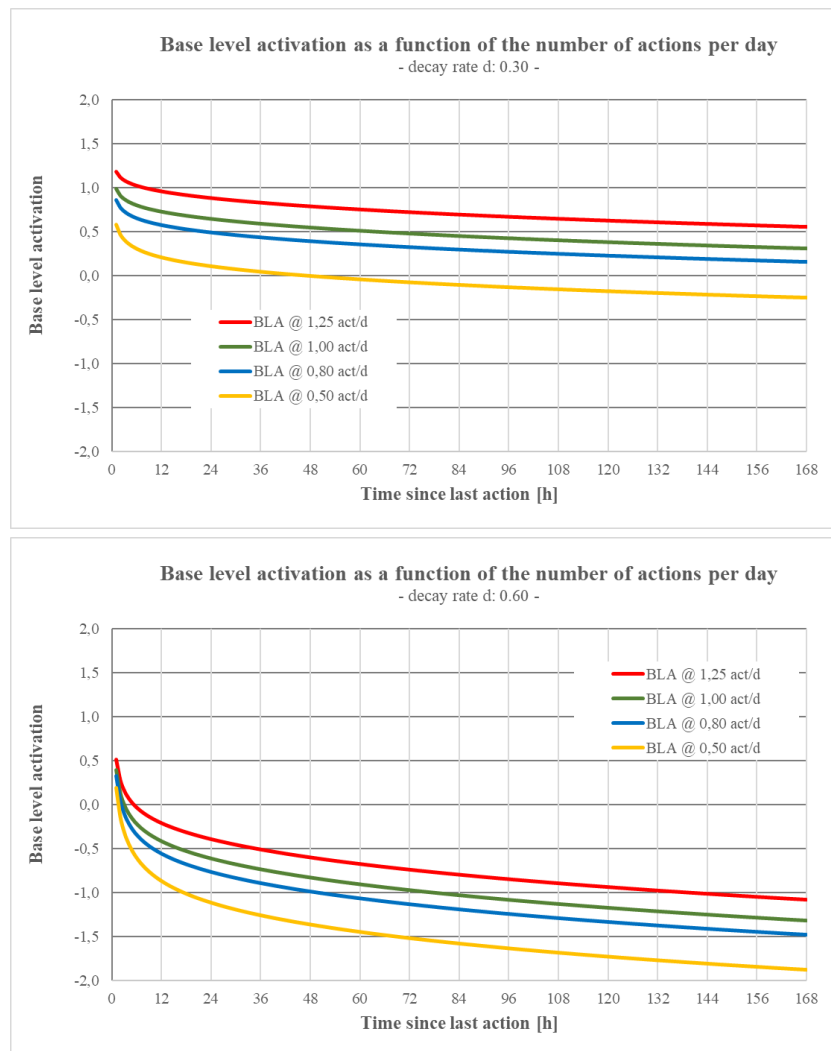


Fig 17: Influence of the decay rate on base-level activation for different, evenly distributed frequencies of interactions ranging from 0.5 actions to 1.25 actions per 24h.

3.5.2 Effect of base-level activation on model performance

Based on the reasoning explicated in the preceding section, a set of decay parameters ranging between 0.30 and 0.60 was used to study the effect of base-level activation on the model performance. Performance improvements cannot be observed, which is in line with the assumption

that base level activation can be considered a process that increases the decision efficiency in a patterned environment but does not necessarily increase the precision of the decision process. As illustrated by Fig 18, a decay value of $d = 0.50$ seems to slightly impair the model performance for thermal sensation in comparison with not using base-level activation; however, this difference is not significant ($p = 0.41$). Both lower and higher decay values tend to decrease the model performance; however, the differences are small, and only the differences for $d = 0.3$ and the version without decay are almost significant ($p = 0.08$). No changes can be observed for air quality sensation; again, this is likely due to the already high performance in this need category.

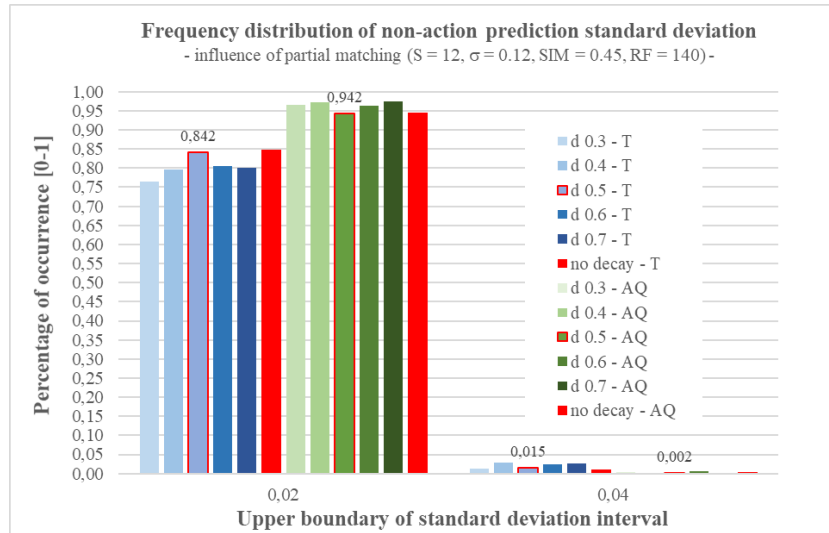


Fig 18: Influence of base-level activation on the model performance.

Analysis of quality of decision did not yield any further information and is therefore not presented here.

4. Discussion and Conclusion

This article presented a study designed to parametrize an instance-based learning and decision model for the simulation and prediction of energy-relevant human-building interaction. The study is part of a larger research endeavor that aims at developing a psychologically based action model as an alternative to the probabilistic approaches currently prevalent in the research field. The line of research, which was roughly outlined in [64], includes 1) the development and application of a qualitative method for the systematic identification and organization of the context of energy-relevant human interaction with buildings [34]; 2) a description of the process of transferring this qualitative data into a numerical format, as exemplified by the development of a numerical decision framework [65]; and 3) the step-by-step advancement of this framework into a cognitively based and psychologically plausible numerical model [28, 38, 43, 66]. Apart from a few exceptions, which are described throughout the text, the model that was parametrized in this study corresponds precisely with the model that was explained in detail in [28]. However, the arithmetic of the model was too extensive to be described coherently and in detail in this article, although parts of the theory were explained when necessary for the parametrization process. Therefore, it is necessary to consult [28] for a better understanding of the details of the model.

The study was separated into the parametrization of five distinct model parts, which were treated step by step. Since a mutual influence between the involved parameters cannot be excluded, this

approach must be considered heuristic, which comes close to an optimum parametrization but is not necessarily the absolute optimum. This approach is an appropriate compromise between accuracy and effort and there is currently no reason to believe that the parametrization was insufficient. Much importance was attached to making a theoretically and empirically informed choice rather than merely choosing an arbitrary set of parameter values for the testing of the model performance. Therefore, each subsection dealing with a specific parameter was separated into a first part encompassing a theoretical introduction, analytical evaluation, and empirical foundation of the parameter choice, as well as a second part elucidating the simulation-based testing of the pre-selected range of parameters. In each of the cases, the final parameter choice was based on a criterion that evaluated the internal consistency of the non-action predictions of the model. This criterion was calculated as the standard deviation between the three non-action predictions related to the three available action dimensions. The more frequent the model produced low standard deviations in a simulation run, i.e., the more consistent the predictions were, the higher the assessments of the model performance were rated. The strength of this criterion is its independence of empirical action data, which makes the parameter choice less specific to a particular situation and thus more generally applicable. In cases when results were ambiguous, additional evaluation criteria were used.

The baseline for the performance-based parameter study was set by establishing the similarity parameter, which is essential for further parametrization because it realistically discretizes the sensation space into bands of similar sensation. The choice of the similarity value was based on findings from two studies in the field of thermal sensation [41] [42]. Both studies used air temperature to test thermal sensation JNDs, which, after transfer of the data into PMV-based thermal sensation values, resulted in a similarity value of 0.45, which was then used for the remainder of the study. However, the substantially lower discriminative abilities of older participants must be considered in simulations in which the building population includes such subjects. Moreover, due to the lack of appropriate data, the similarity value was assumed to be valid across modalities and was therefore also used for air quality sensation. Future studies will need to test if this assumption is valid or if modality-specific values must be established.

The second treated parameter was the maximum associative strength, which is relevant for the determination of the contextually cued retrieval process of information from the declarative memory. Previous versions of the model demonstrated that a value of around two, which is the common choice for this parameter in cognitive models, is inappropriate for simulations of energy-relevant human action in buildings. Based on the plausibility of the results, these models rather used a value of 30. Based on the number of instances built-in declarative memory and the corresponding “fan,” this study demonstrated that a maximum associative strength of around 12 is appropriate and showed that values below this figure substantially degraded the model performance.

Noise adds an element of indeterminacy to the retrieval process and triggers the model to occasionally choose seemingly non-optimal actions; however, this also leads to a positive further exploration of the world and thus increases the experience and abilities of the model. Nonetheless, this effect is limited, as too-high noise values dilute the associative retrieval process and thus randomizes the decisions of the model. The parametrization process adjusted the noise value for the model to 0.12 to achieve a balance between its exploration of the world and its exploitation of known solutions to pre-learned problems.

The introduction and parametrization of partial matching improved the model performance substantially by providing it with the ability to differentiate the degree of mismatch between environmental cues and declaratively stored information. Partial matching penalizes activation in proportion to the degree of mismatch, such that substantially mismatching slots can render the total activation of that instance to be zero or even negative even though the remaining slots fit the context. This is particularly relevant for the model performance because a contextual mismatch for the action state of an element can reverse the effect of the action. To simplify the handling of partial matching, slot-specific partial matching values were set in such a manner that multiplication of the partial matching value with the maximal state difference yielded practically identical maximal penalization values for each slot. As a consequence, all slots were treated as equally important and as yet, there has been no indication that this was an inappropriate assumption. The best model performance turned out to be yielded by a maximal penalization of six, which corresponds well with the maximum associative strength S of 12 and the typical $\ln(fan)$ of six.

Finally, the concept of information decay was introduced into the model through the utilization of base-level activation. By analyzing empirical data, the decay value was theoretically limited to a value between 0.3 and 0.6 and finally set to 0.5 in alignment with previous versions of the model.

The parameter optimization process illustrated in the article led to a substantial improvement of the model performance as measured by the internal predictive consistency. This is an important step towards a realistic and psychologically based model of human behavior in buildings. Nonetheless, as stated throughout the article, the model is restricted to certain contexts and still lacks some human abilities. The most important of these restrictions is certainly the lack of an occupancy model, resulting in the model to unrealistically act 24h per day. Future work thus needs to address this lack by integrating a context-specific occupancy model. However, closely related to the question of occupancy is the ability of the model to plan over periods of absence (such as overnight ventilation, work times away from home). The current decision horizon of the model covers one hour which is inappropriate for extended leaves of the room. Consequently, the psychological mechanisms involved in long-term planning need to be developed and integrated in the model in the future. Furthermore, the model is currently restricted to a setting with single occupancy. Future enhancement of the model thus also needs to integrate social aspects in the decision process which occur when people share spaces. Finally, the model results need to be tested against empirical behavior data to test the validity of the model. Even though a fully developed model would be preferable to compare individual behavior against model results, a paper currently in preparation will demonstrate that even at the current developmental stage the model is already able to explain human behavior at the aggregate level.

References:

1. Clarke, J.A., I. Macdonald, and J.F. Nicol, *Predicting adaptive responses-simulating occupied environments*, in *In: Proceedings of International Conference on Comfort and Energy Use in Buildings*. 2006: Windsor, England, .
2. Haldi, F. and D. Robinson, *On the behaviour and adaptation of office occupants*. Building and Environment, 2008. **43**(12): p. 2163-2177.
3. Haldi, F. and D. Robinson, *Interactions with window openings by office occupants*. Building and Environment, 2009. **44**(12): p. 2378-2395.
4. Haldi, F. and D. Robinson, *Modelling occupants' personal characteristics for thermal comfort prediction*. International journal of biometeorology, 2011. **55**(5): p. 681-694.
5. Inkarojrit, V. and G. Paliaga. *Indoor climatic influences on the operation of windows in a naturally ventilated building*. in *Proceedings of the 21st international conference on passive and low energy architecture*. 2004. Eindhoven, The Netherlands.
6. Li, N., et al., *Probability of occupant operation of windows during transition seasons in office buildings*. Renewable Energy, 2015. **73**: p. 84-91.
7. Newsham, G.R., *Manual control of window blinds and electric lighting: implications for comfort and energy consumption*. Indoor and Built Environment, 1994. **3**(3): p. 135-144.
8. Nicol, J.F., *Characterising occupant behaviour in buildings: towards a stochastic model of occupant use of windows, lights, blinds, heaters and fans*, in *Seventh International IBPSA Conference*. 2001: Rio de Janeiro. p. 1073-1078.
9. Ren, X., D. Yan, and C. Wang, *Air-conditioning usage conditional probability model for residential buildings*. Building and Environment, 2014. **81**: p. 172-182.
10. Rijal, H.B., et al., *Considering the impact of situation-specific motivations and constraints in the design of naturally ventilated and hybrid buildings*. Architectural Science Review, 2012. **55**(1): p. 35-48.
11. Rijal, H.B., et al., *Using results from field surveys to predict the effect of open windows on thermal comfort and energy use in buildings*. Energy and buildings, 2007. **39**(7): p. 823-836.
12. Rijal, H.B., et al., *Development of adaptive algorithms for the operation of windows, fans and doors to predict thermal comfort and energy use in Pakistani buildings*. ASHRAE transactions, 2008. **114**(2): p. 555-573.
13. Wang, C., et al., *A generalized probabilistic formula relating occupant behavior to environmental conditions*. Building and Environment, 2016. **95**: p. 53-62.
14. Warren, P.R. and L.M. Parkins, *Window-opening behaviour in office buildings*. Building Services Engineering Research and Technology, 1984. **5**(3): p. 89-101.
15. Widén, J., A.M. Nilsson, and E. Wäckelgård, *A combined Markov-chain and bottom-up approach to modelling of domestic lighting demand*. Energy and Buildings, 2009. **41**(10): p. 1001-1012.
16. Zhang, Y. and P. Barrett, *Factors influencing occupants' blind-control behaviour in a naturally ventilated office building*. Building and Environment, 2012. **54**: p. 137-147.
17. Foster, M. and T. Oreszczyn, *Occupant control of passive systems: the use of venetian blinds*. Building and Environment, 2001. **36**(2): p. 149-155.
18. Zhang, Y. and P. Barrett, *Factors influencing the occupants' window opening behaviour in a naturally ventilated office building*. Building and Environment, 2012. **50**: p. 125-134.
19. Gunay, H.B., W. O'Brien, and I. Beausoleil-Morrison, *A critical review of observation studies, modeling, and simulation of adaptive occupant behaviors in offices*. Building and Environment, 2013. **70**: p. 31-47.
20. Stazi, F. and F. Naspi, *Modelling, Implementation and Validation Approaches*, in *Impact of Occupants' Behaviour on Zero-Energy Buildings*. 2018, Springer International Publishing: Cham. p. 63-77.
21. Kim, D.-W., et al. *Traditional vs. cognitive agent simulation*. in *13th International conference of the international building performance simulation association*. 2013.
22. Langevin, J., J. Wen, and P.L. Gurian, *Simulating the human-building interaction: Development and validation of an agent-based model of office occupant behaviors*. Building and Environment, 2015. **88**: p. 27-45.

23. Lee, Y.S. and A.M. Malkawi, *Simulating multiple occupant behaviors in buildings: An agent-based modeling approach*. Energy and Buildings, 2014. **69**: p. 407-416.
24. Hong, T., et al., *An occupant behavior modeling tool for co-simulation*. Energy and Buildings, 2016. **117**: p. 272-281.
25. Hong, T., et al., *Advances in research and applications of energy-related occupant behavior in buildings*. Energy and Buildings, 2016. **116**: p. 694-702.
26. Hong, T., et al., *Ten questions concerning occupant behavior in buildings: The big picture*. Building and Environment, 2017. **114**: p. 518-530.
27. Yan, D., et al., *Occupant behavior modeling for building performance simulation: Current state and future challenges*. Energy and Buildings, 2015. **107**: p. 264-278.
28. Grabe, J.v., *Using the instance-based learning paradigm to model energy-relevant occupant behaviors in buildings* Cognitive Computation, 2020. **12**: p. 71-99.
29. Funke, J., *Problemlösendes Denken*. 2003: Kohlhammer Verlag.
30. Anderson, J.R., et al., *An integrated theory of the mind*. Psychological Review, 2004. **111**(4): p. 1036-1060.
31. Gonzalez, C., J.F. Lerch, and C. Lebiere, *Instance-based learning in dynamic decision making*. Cognitive Science, 2003. **27**(4): p. 591-635.
32. Taatgen, N.A., C. Lebiere, and J.R. Anderson, *Modeling paradigms in ACT-R*, in *Cognition and multi-agent interaction: From cognitive modeling to social simulation*, R. Sun, Editor. 2006, Cambridge University Press. p. 29-52.
33. Korkas, C.D., et al., *Intelligent energy and thermal comfort management in grid-connected microgrids with heterogeneous occupancy schedule*. Applied Energy, 2015. **149**: p. 194-203.
34. Grabe, J.v., *The systematic identification and organization of the context of energy-relevant human interaction with buildings—a pilot study in Germany*. Energy Research & Social Science, 2016. **12**: p. 75-95.
35. Fanger, P.O., *Thermal comfort: analysis and applications in environmental engineering*. 1970, Copenhagen: Danish Technical Press.
36. Fanger, P.O., *Introduction of the olf and the decipol units to quantify air pollution perceived by humans indoors and outdoors*. Energy and Buildings, 1988. **12**(1): p. 1-6.
37. Gunnarsen, L. and P.O. Fanger, *Adaptation to indoor air pollution*. Environment International, 1992. **18**(1): p. 43-54.
38. Grabe, J.v., *A preliminary cognitive model for the prediction of energy-relevant human interaction with buildings*. Cognitive Systems Research, 2018. **49**: p. 65-82.
39. Klein, A., et al. *TRNSYS 17—a transient system simulation program*. 2010 4/11/2017]; Available from: <http://sel.me.wisc.edu/trnsys>.
40. Fechner, G.T., *Elemente der Psychophysik - Erster Theil*. Vol. 1. 1860, Leipzig: Druck und Verlag von Breitkopf und Härtel.
41. Collins, K.J., A.N. Exton-Smith, and C. Doré, *Urban hypothermia: preferred temperature and thermal perception in old age*. Br Med J (Clin Res Ed), 1981. **282**(6259): p. 175-177.
42. Natsume, K., et al., *Preferred ambient temperature for old and young men in summer and winter*. International journal of biometeorology, 1992. **36**(1): p. 1-4.
43. Grabe, J.v. and C. Gonzalez, *Human decision making in energy-relevant interaction with buildings*, in *Central European Symposium on Building Physics, CESBP*. 2016: Dresden. p. 345-352.
44. Anderson, J.R., *A spreading activation theory of memory*. Journal of Verbal Learning and Verbal Behavior, 1983. **22**(3): p. 261-295.
45. Schooler, L.J. and J.R. Anderson, *Recency and Context: An Environmental Analysis of Memory*, in *Fifteenth Annual Conference of the Cognitive Science Society*. 1993. p. 889-894.
46. Anderson, J.R., *Retrieval of propositional information from long-term memory*. Cognitive psychology, 1974. **6**(4): p. 451-474.
47. Anderson, J.R. and L.M. Reder, *The fan effect: New results and new theories*. Journal of Experimental Psychology-General, 1999. **128**(2): p. 186-197.
48. Thomson, R. and C. Lebiere. *A balanced Hebbian algorithm for associative learning in ACT-R*. in *Proceedings of the International Conference on Cognitive Modeling*. 2013.

49. Lejarraga, T., V. Dutt, and C. Gonzalez, *Instance-based learning: A general model of repeated binary choice*. Journal of Behavioral Decision Making, 2010. **25**(2): p. 143-153.
50. Gonzalez, C. and V. Dutt, *Instance-based learning: Integrating sampling and repeated decisions from experience*. Psychological review, 2011. **118**(4): p. 523.
51. Dutt, V. and C. Gonzalez, *Making instance-based learning theory usable and understandable: The Instance-Based Learning Tool*. Computers in Human Behavior, 2012. **28**(4): p. 1227-1240.
52. Reder, L.M. and G.W. Kusbit, *Locus of the Moses illusion: Imperfect encoding, retrieval, or match?* Journal of Memory and Language, 1991. **30**(4): p. 385-406.
53. Erickson, T.D. and M.E. Mattson, *From words to meaning: A semantic illusion*. Journal of Verbal Learning and Verbal Behavior, 1981. **20**(5): p. 540-551.
54. Berry, D.C. and D.E. Broadbent, *On the relationship between task performance and associated verbalizable knowledge*. The Quarterly Journal of Experimental Psychology, 1984. **36**(2): p. 209-231.
55. Taatgen, N.A. and D. Wallach, *Whether skill acquisition is rule or instance based is determined by the structure of the task*. Cognitive Science Quarterly, 2002. **2**(2): p. 163-204.
56. Siegler, R.S. and J. Shrager, *Strategy choices in addition and subtraction: How do children know what to do*. Origins of cognitive skills, 1984. **23**(1): p. 229-293.
57. Anderson, J.R. and L.J. Schooler, *Reflections of the environment in memory*. Psychological science, 1991. **2**(6): p. 396-408.
58. Rakow, T. and K. Miler, *Doomed to Repeat the Successes of the Past: History Is Best Forgotten for Repeated Choices with Nonstationary Payoffs*. Memory & Cognition, 2009. **37**(7): p. 985.
59. Anderson, J.R., J.M. Fincham, and S. Douglass, *Practice and retention: A unifying analysis*. Journal of Experimental Psychology-Learning Memory and Cognition, 1999. **25**(5): p. 1120-1136.
60. Pavlik, P.I. and J.R. Anderson, *An ACT-R model of the spacing effect*. Department of Psychology, 2003: p. 56.
61. Pavlik, P.I. and J.R. Anderson, *Practice and forgetting effects on vocabulary memory: An activation-based model of the spacing effect*. Cognitive Science, 2005. **29**(4): p. 559-586.
62. Morgan, C. and R. de Dear, *Weather, clothing and thermal adaptation to indoor climate*. Climate Research, 2003. **24**(3): p. 267-284.
63. de Carvalho, P.M., M.G. da Silva, and J.E. Ramos, *Influence of weather and indoor climate on clothing of occupants in naturally ventilated school buildings*. Building and environment, 2013. **59**: p. 38-46.
64. Grabe, J.v., *A psychological approach to understanding and predicting energy-relevant human interaction with buildings*, in *Passive and Low Energy Architecture, PLEA*. 2016: Los Angeles.
65. Grabe, J.v., *How do occupants decide their interactions with the building? From qualitative data to a psychological framework of human-building-interaction*. Energy Research & Social Sciences, 2016. **14**: p. 46-60.
66. Grabe, J.v., *Cognitive Modelling for the Prediction of energy-relevant Human Interaction with Buildings*, in *13th International Conference on Cognitive Modelling*, N.A. Taatgen, et al., Editors. 2015: Groningen.