TU Delft

Delft University of Technology

Document Version
Final published version

Licence
Dutch Copyright Act (Article 25fa)

Citation (APA)
Martinez-Alcaraz, P., Barra, P. D. L., Andriotis, C. P., & Luna-Navarro, A. (2026). Does One Size Fit All? A Bayesian Data-Mining Approach to Thermal Comfort in Office Spaces. In C. Zilio, F. Busato, L. Mazzarella, & M. Noro (Eds.), *Proceedings of the 15th REHVA HVAC World Congress - CLIMA 2025* (Vol. 3, pp. 1454-1463). (Lecture Notes in Civil Engineering; Vol. 788 LNCE). Springer. https://doi.org/10.1007/978-3-032-10546-2_138

Important note
To cite this publication, please use the final published version (if applicable).
Please check the document version above.

Copyright
In case the licence states “Dutch Copyright Act (Article 25fa)”, this publication was made available Green Open Access via the TU Delft Institutional Repository pursuant to Dutch Copyright Act (Article 25fa, the Taverne amendment). This provision does not affect copyright ownership.
Unless copyright is transferred by contract or statute, it remains with the copyright holder.

Sharing and reuse
Other than for strictly personal use, it is not permitted to download, forward or distribute the text or part of it, without the consent of the author(s) and/or copyright holder(s), unless the work is under an open content license such as Creative Commons.

Takedown policy
Please contact us and provide details if you believe this document breaches copyrights.
We will remove access to the work immediately and investigate your claim.

This work is downloaded from Delft University of Technology.

Does One Size Fit All? A Bayesian Data-Mining Approach to Thermal Comfort in Office Spaces

P. Martinez-Alcaraz(✉), P. de la Barra, C. P. Andriotis, and A. Luna-Navarro

Department of Architectural Engineering and Technology, Delft University of Technology, Delft, The Netherlands
p.martinezalcaraz@tudelft.nl

Abstract. Building Automation and Control Systems (BACS) enhance energy efficiency in office buildings but often leave occupants dissatisfied, especially in shared offices with diverse indoor environmental preferences. Personal Comfort Models (PCM) capture individual preferences, yet these models can be difficult to apply in shared environments. This study examines whether a generalized personalized modeling would be necessary, or instead aggregated thermal preferences adequately represent individual needs in shared spaces. By mining the ASHRAE Global Thermal Comfort Database II, we analyzed 11 naturally ventilated office buildings spanning four climate zones, each providing 187 to 568 data points of subjective thermal preferences and concurrent environmental measurements. To capture uncertainties in occupant feedback and monitoring, we developed a Bayesian multinomial logistic framework. We compare an aggregated model (all responses together) with a hierarchical (partial-pooling) model having occupant-specific intercepts, using the No-U-Turn Sampler for Hamiltonian Monte Carlo. The Widely Applicable Bayesian Information Criterion confirms that hierarchical approach outperforms the aggregated model in every case. In four of the eleven studies, aggregated air temperature ranges satisfied occupants' demands within a 95% posterior credible interval. Moreover, in most buildings, these temperature ranges aligned well with the majority of occupant preferences, suggesting that prior knowledge of building conditions or occupant variability might reduce the need for highly customized thermal environment. Future research should explore these findings in conditioned buildings and across other indoor environmental quality domains, where greater variability could exist.

Keywords: Thermal Comfort · Occupant-Centric Building Controls · Bayesian · Data-Mining · Personal Comfort Model

Nomenclature

Acronyms

BACS	Building Automation and Control Systems
PCM	Personal Comfort Models
HVAC	Heating, Ventilation and Air-Conditioning
ASHRAE	American Society of Heating, Refrigerating and Air-Conditioning Engineers

© The Author(s), under exclusive license to Springer Nature Switzerland AG 2026
C. Zilio et al. (Eds.): CLIMA 2025, LNCE 788, pp. 1454–1463, 2026.
https://doi.org/10.1007/978-3-032-10546-2_138

NUTS	No-U-Turn Sampler
HMC	Hamiltonian Monte Carlo
WBIC	Widely Applicable Bayesian Information Criterion
HDI	Highest Density Intervals

1 Introduction

Building Automation and Control Systems (BACS) are instrumental in improving the energy efficiency and operational performance of non-residential buildings by coordinating services such as HVAC, lighting, and shading devices [1]. Often, BACS are seen as disruptive for occupants, due to minimal personal control over indoor environmental conditions and control strategies that fail to account for occupant requirements in real time [2, 3].

Occupant-Centric Controls (OCC) place the human at the core of building management, addressing dissatisfaction by dynamically adjusting environmental settings based on real-time occupant feedback, which can boost Indoor Environmental Quality (IEQ) and energy performance [4]. Research initiatives such as the IEA EBC Annex 79 highlight human-building interaction as a way to enhance overall building performance [5]. Within this OCC framework, Personal Comfort Models (PCMs) have gained prominence for capturing individual occupant behaviors more granularly based on subjective feedback (e.g., thermal sensation votes) and environmental monitoring (e.g., temperature, humidity) [6].

Extending PCMs to shared office environments is challenging. Occupant-centric strategies must identify who is present and how they prefer the space. This typically calls for extra technological infrastructure, such as digital IDs or occupant detection, which is not yet common in many offices. In addition, occupants often provide subjective feedback only when uncomfortable, leading to unbalanced datasets and biased inference.

Research in thermal comfort has shown that broad temperature ranges can simultaneously meet the preferences of several occupants [7], while personal comfort models (PCM) aim at providing tailored conditions that can meet individual preferences. The objective of this research is to determine when occupant-specific comfort models are necessary and when they are not required. In other words, *"Is a universal approach sufficient, or do we need personalized solutions in all cases?"*. By comparing both aggregated and occupant-level comfort models, this study discusses under which conditions personal comfort strategies could provide higher occupant satisfaction.

To this end, we applied a data-mining approach on selected 11 office building studies from the ASHRAE Global Thermal Comfort Database II. We developed two models: an aggregated one that considers all occupant responses as a single group, and a PCM that captures individual preferences. Since occupant preferences are inherently uncertain, we adopt probabilistic methods to reflect the stochastic nature of IEQ and occupant feedback [9].

2 Methodology

We propose a data-mining approach that combines occupant comfort data from the ASHRAE Global Thermal Comfort Database II. We use a multinomial logistic framework to predict three thermal preferences: "cooler", "no change", and "warmer". We compare two Bayesian modeling strategies: an aggregated model and a hierarchical model that includes occupant-specific intercepts. Model inference relies on the No-U-Turn Sampler (NUTS) for Hamiltonian Monte Carlo. The Widely Applicable Bayesian Information Criterion (WBIC) then guides models posteriors understanding. We generate posterior predictive probabilities over a range of indoor temperatures and hold other predictors at their mean values to determine whether aggregated temperature ranges would be enough to represent the majority of occupants' preferences.

2.1 Data and Thermal Comfort Definition

We used the ASHRAE Global Thermal Comfort Database II [10], which aggregates data from various buildings, occupants, and environmental conditions worldwide. We selected 11 buildings in which each occupant had at least ten entries covering indoor air temperature (°C), indoor relative humidity (%), indoor air velocity (m/s), outdoor air temperature (°C), outdoor relative humidity (%), and thermal preference responses (t_{pref}). These t_{pref} responses were subjective categories such as "warmer", "no change", or "cooler". All selected studies were conducted in naturally ventilated office buildings across four climate zones (desert, semi-arid high altitude, semi-arid midlatitude, and Mediterranean) located in Asia and Europe. Of the 11 studies, eight took place in summer (IDs 172, 174, 177, 179, 175, 141, 143, and 145) and three in winter (IDs 173, 178, and 176). Each study involved three to six occupants, yielding data totals between 187 and 568 points per study (about 15 to 117 points per occupant).

For predicting thermal comfort, a data-driven thermal preference model based on a Multinomial Logistic Regression that uses the SoftMax function is learned (Lee et al., 2017; Ma et al., 2021). This function estimates the probability that an observation i belongs to one of K classes (e.g., "cooler", "no change", or "warmer"), based on its input features $x^{(i)} = \{x_1^{(i)}, x_2^{(i)}, \ldots, x_M^{(i)}\}$ (e.g., indoor air temperature, indoor relative humidity, outdoor temperature, and outdoor relative humidity). In this model, w is a vector of weights, b represents the intercept value, i is a training example, l is the feature index, and M is the number of features:

$$P\left(t_{pref} = k \middle| z^{(i)}\right) = \frac{e^{Z^{(i)}}}{\sum_{k=1}^{K} e^{Z_K^{(i)}}}, z = b + \sum_{L=0}^{M} w_L x_L \tag{1}$$

2.2 Aggregated and Individual Bayesian Models

We applied two Bayesian multinomial logistic models to each study. The aggregated model treated all occupant data as if from one group. The hierarchical (partial-pooling)

model assigned occupant-specific intercepts drawn from hierarchical *Normal* and *Half-Cauchy* hyperpriors. The *Normal* hyperprior centers individual intercepts around a common group mean, while the *HalfCauchy* hyperprior provides a robust, weakly informative estimate of variability, allowing the model to borrow strength from the entire group, particularly for occupants with sparse or outlier data. Both models used weakly informative *Normal (0, 1)* priors for their global coefficients and intercepts.

We estimated parameters in Python's Pyro using NUTS. We ran four Hamiltonian Monte Carlo chains, each with 500 warm-up steps and 1,000 sampling steps. We checked convergence via standard diagnostics and extracted 95% credible intervals from the posterior samples.

2.3 Analysis and Comparison

To assess the benefit of accounting for occupant-specific variation, we compare two Bayesian modeling approaches: an aggregated model treating all responses as a single group and a hierarchical model with occupant-specific intercepts. Because our primary goal is to compare the full posterior distributions of these models, we use the Widely Applicable Bayesian Information Criterion (WBIC) for model selection. WBIC approximates the log marginal likelihood by using a tempered posterior and then scaling the average total log likelihood accordingly. Lower WBIC values indicate a model with higher evidence, reflecting a better trade-off between fit and complexity. By comparing the WBIC scores of the aggregated and hierarchical models, we determine whether incorporating occupant-specific parameters yields a more informative posterior, thereby justifying the added complexity.

Although our model includes humidity, air velocity, and outdoor conditions, we focused on air temperature because it is both the most directly controlled variable and highly relevant to occupant comfort and energy efficiency. First, we assessed the significance of each predictor by examining whether its highest posterior density interval (HDI) excluded zero. Next, to understand air temperature ranges, we computed posterior predictive probabilities over a 100-point grid of indoor air temperature values while holding the other variables at their mean levels. For each temperature, we then derived the 95% credible interval for the predicted probability of each thermal response class (i.e., "cooler", "no change", and "warmer"). Subsequently, we defined the comfort range as the set of temperatures where the 95% credible interval for the "no change" response indicates at least a 95% posterior probability of "no change" being the most likely outcome—exceeding the probabilities for both the "cooler" and "warmer" classes. Thereafter, we calculated the number of occupants whose "no change" temperature range, as determined by the 95% credible interval, matched that of the aggregated model. When an occupant's "no change" range was either unresolved or spanned multiple categories, we marked that range as missing, indicating insufficient data to define distinct comfort limits. Finally, We computed the Pearson correlation coefficient between the aggregated and hierarchical model predictions for the 'no change' class. Correlation coefficient above 0.90 represent a very strong correspondence, indicating that both models yield nearly identical predicted probabilities across the range of temperatures. In this way, by examining the aggregated and occupant-specific outcomes together, we determined whether

a single temperature range could reliably serve most occupants or whether variability in individual preferences required a more personalized modeling approach.

3 Results

We begin by describing the results from our modeling approach for Study 143. This example highlights how occupant-level variation in temperature preferences can lead to different predictive outcomes compared to an aggregated approach. We then extend our analysis to all 11 studies to see whether incorporating individual granularity would be relevant.

3.1 Model Comparison and Posterior Predictions: Study 143 as Example

Study 143 features average indoor relative humidity of 31%, air velocity of 0.71 m/s, outdoor air temperature of 23.5 °C, and outdoor relative humidity of 56%. In the aggregated model (Fig. 1), these predictors are held constant while indoor air temperature varies from 20.8 °C to 31 °C. The "no change" preference dominates across much of this range, peaking near 23.4 °C, and its 95% credible interval remains from ~20.8 °C to ~31 °C. This means that at temperatures within this interval, there is a 95% probability, according to the posterior distribution, that the "no change" response is the most likely outcome, thereby defining the occupant's comfort range. In addition, WBIC results indicates that the aggregated model have better presentation (WBIC = 1292) compared to the hierarchical model (WBIC = 1322).

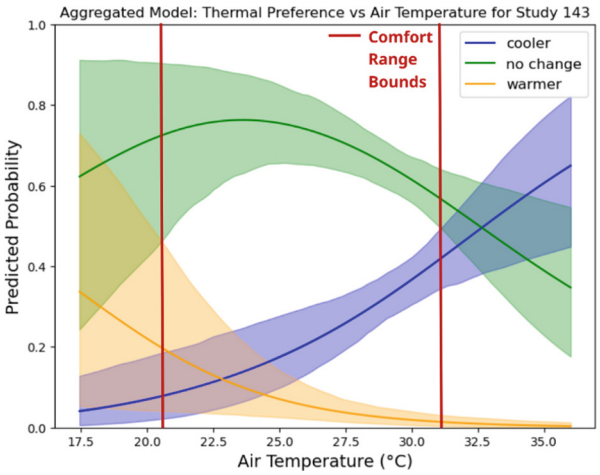


Fig. 1. Aggregated model posterior prediction of thermal preference over air temperature while keeping the other predictors steady. In red, the temperature range bounds where 95% credible interval for class no change is greater than others (Study 143).

Correlations between hierarchical and aggregated predictions also reveal individual differences in Study 143. Three occupants exhibit very strong correlation coefficients (0.9–1.0), while the other three show moderate or low correlation coefficients (< 0.70). Consequently, 50% of occupants have air temperature range that fall outside the aggregated model's range. Figure 2 illustrates these distinctions. For example, Occupants 143_14, 143_15, and 143_16 show overlapping probabilities for multiple preference classes with not a clear preferred air temperature range, whereas Occupants 143_17 and 143_19 have stronger "no change" peaks (~23.6 °C–30.8 °C and ~24.3 °C–32.8 °C, respectively). Occupant 143_18 demonstrates the widest preferred air temperature range (~17.5 °C–32.25 °C). These variations highlight how partial pooling can capture occupant-level nuances that an aggregated model overlooks.

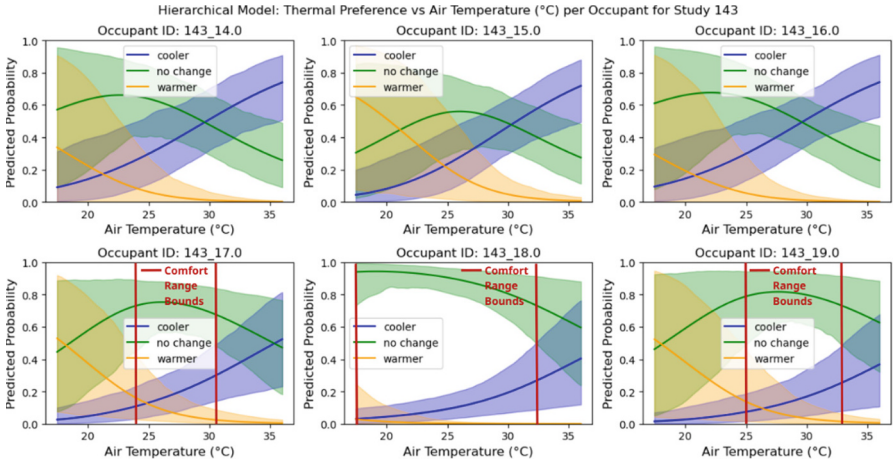


Fig. 2. Hierarchical Model posteriors of thermal preference over air temperature while keeping the other predictors steady. In red, the temperature range bounds where 95% credible interval for class no change is greater than others. Occupants ID 143_14, 143_15 and 143_16 do not show preferred comfort range with 95% credible interval (Occupant ID for Study 143).

3.2 Models Comparison and Posterior Predictions

Table 1 presents the WBIC metrics for all 11 studies. In 10 out of the 11 studies, the hierarchical model has a lower (i.e., better) WBIC than the aggregated model, often by a substantial margin (e.g., 161 points in Study 172 or 604 points in Study 176). In one case (Study 143), the difference is negative ($\Delta\text{WBIC} = -30$), indicating that the aggregated model has slightly higher evidence than the hierarchical approach.

Table 1. WBIC, and number of data points per Study.

Study ID	Num. Data points	WBIC Agg. Model	WBIC Hier. Model	ΔWBIC (Agg – Hier)
172	190	832	671	161
173	470	2897	2537	361
174	436	1728	1598	130
177	491	2173	2157	16
178	424	1893	1756	138
179	568	2056	1626	430
175	547	2975	2522	452
176	511	3061	2457	604
141	280	1028	867	161
143	279	1292	1322	−30
145	187	613	566	47

Using a Bayesian data-mining approach across multiple office buildings, we computed 95% highest density intervals (HDIs) for the model parameters. Notably, only air temperature ("ta") showed a credible interval that excludes zero for both cooler and warmer outcomes, indicating a significant effect on occupant thermal preferences. By contrast, the effects of other variables did not credibly deviate from zero, suggesting their influence on thermal comfort is marginal or negligible in those contexts. The correlation coefficients analysis (Table 2) reveals that, while global predictions often align with individual patterns (correlation coefficients > 0.95 in many instances), there are some exceptions. For example, Studies 178, 143 and 175 display relatively low correlations coefficients for some occupants in comparison with the aggregated model, further illustrating that global trends can fail to capture individual preferences. When examining the air temperature range (lower, optimal, and upper) for the "no change" category, similar patterns emerge. In studies such as 174, 177, 178 and 145, nearly all occupants' range align closely with the aggregated estimates, suggesting that preferred air temperature based on the aggregated model might comfortably include the majority of occupants demand. By contrast, in studies 172, 173, 179, 175, 176, 141, and 143, a small proportion of occupants have unresolved ranges in the model, which can be due to the lack of data dispersion, or ranges that differ substantially from the aggregated model. In the Study 141, for instance, the aggregated model presents unresolved ranges. There, 86% of occupants might fall out of aggregated air temperature ranges.

Table 2. Correlation coefficients of hierarchical models versus the aggregated one, number of occupants and percentage of occupants outside of the air temperature ranges of the aggregated model.

Study ID	Mean rh (%)	Mean airV (m/s)	Range ta [Lower, Upper]	N. of Occ	% of Occ. Outside ta Ranges	Very Strong Corr. Coef (0.9–1.0)	Strong Corr. Coef (0.7–0.89)	Mod./Low Corr. Coef. (<0.70)
172	65	0.39	[29, 31]	4	50	3	1	0
173	42	0.09	[23.6, 29.8]	5	40	5	0	0
174	69	0.29	[24.6, 29.1]	5	0	4	1	0
177	23	0.30	[18.3, 27.5]	6	0	5	1	0
178	28	0.18	[13.91, 25.8]	5	0	1	2	2
179	46	0.19	[19.5, 27.2]	5	20	2	3	0
175	51	0.24	[24.2, 31.1]	6	33	3	2	1
176	57	0.15	[18.2, 31.4]	5	40	4	1	0
141	39	0.35	[-, -]*	7	86*	5	1	1
143	31	0.71	[20.8, 31]	6	50	3	0	3
145	30	0.32	[23, 30.6]	3	0	3	0	0

* Unresolved model prediction: the model never reaches a 95% posterior probability for the "no change" class, so no temperature range can be established.

4 Discussion

Overall, these results suggest that modeling individual differences with per-occupant intercepts significantly improves model evidence for most buildings, though occupant-level modeling may not always lead to better representation in every context. A subset of occupants showed air temperature preferences that diverged significantly from the group average, underscoring the potential of more nuanced, individual approaches. Nevertheless, aggregated modeling, which applies a global air temperature range, adequately represented occupant preferences in four out of the 11 examined studies. Furthermore, aggregated temperature ranges encompassed the majority of occupant demands in all but one study. This finding aligns well with adaptive thermal comfort principles observed in naturally ventilated buildings, where occupants habitually adjust clothing, operate

windows, and modify their environment to maintain comfort across a relatively broad temperature range.

The current analysis, focusing exclusively on naturally ventilated office buildings, suggests caution when extending these findings to mechanically ventilated or conditioned building typologies. Still, aggregated models provide a viable initial strategy, particularly for building managers aiming to implement practical, cost-effective solutions that enhance occupant satisfaction with indoor environmental quality.

5 Conclusions

This work investigated whether occupant-specific temperature ranges are justified in shared office buildings or if an aggregated approach can suffice. We analyzed data from the ASHRAE Global Thermal Comfort Database II (11 buildings) using a multinomial logistic framework and compared aggregated versus hierarchical (partial-pooling) Bayesian models.

The hierarchical model, which accounts for occupant-level variation, achieved a consistently better posteriors evidence. Nevertheless, in 4 of the 11 studies, an aggregated approach still adequately matched most occupants' air temperature preferences. Thus, while occupant-specific modeling can better accommodate individual outliers, a single broad range may be sufficient in many buildings. Therefore, this research highlight the use of personal environmental control systems such as personal heaters or fans, when the variance of thermal preferences in the shared spaces cannot be addressed by a centralized HVAC system.

Future studies could apply the hierarchical approach in real-time building control systems to see whether the complexity of occupant-specific modeling translates into an increase acceptance of building controls. Another avenue is to investigate whether other aspects of indoor environmental quality, such as daylight, lighting, acoustics, and air quality, exhibit greater variability across occupants. Ultimately, these research efforts will guide building managers and engineers in deciding when occupant diversity in preferences justifies the added complexity of personalized control strategies.

Acknoledgements. This study has received funding from the European Union under the Horizon Europe Research & Innovation Programme (Grant Agreement no. 101103956 SMARTeeSTORY).

References

1. Tamas, R., Ouf, M.M., O'Brien, W.: A field study on the effect of building automation on perceived comfort and control in institutional buildings. Archit. Sci. Rev. **63**, 74–86 (2020). https://doi.org/10.1080/00038628.2019.1695573
2. Karmann, C., Schiavon, S., Arens, E.: Percentage of commercial buildings showing at least 80% occupant satisfied with their thermal comfort (2018)
3. de la Barra, P,. Martinez-Alcaraz, P., Luna-Navarro, A.: Identification of factors influencing satisfaction with interaction strategies by clustering occupants in buildings. In: Berardi, U. (ed.) Multiphysics and Multiscale Building Physics, pp. 211–220. Springer, Singapore (2025)

4. Nagy, Z., Gunay, B., Miller, C., et al.: Ten questions concerning occupant-centric control and operations. Build. Environ. **242**, 110518 (2023). https://doi.org/10.1016/J.BUILDENV.2023.110518
5. O'Brien, W., Wagner, A., Schweiker, M., et al.: Introducing IEA EBC annex 79: key challenges and opportunities in the field of occupant-centric building design and operation. Build. Environ. **178**, 106738 (2020). https://doi.org/10.1016/J.BUILDENV.2020.106738
6. Kim, J., Schiavon, S., Brager, G.: Personal comfort models – a new paradigm in thermal comfort for occupant-centric environmental control. Build. Environ. **132**, 114–124 (2018). https://doi.org/10.1016/j.buildenv.2018.01.023
7. Arens, E., Humphreys, M.A., de Dear, R., Zhang, H.: Are 'class A' temperature requirements realistic or desirable? Build. Environ. **45**, 4–10 (2010). https://doi.org/10.1016/J.BUILDENV.2009.03.014
8. Nicol, J.F., Humphreys, M.A.: Adaptive thermal comfort and sustainable thermal standards for buildings. Energy Build. **34**, 563–572 (2002). https://doi.org/10.1016/S0378-7788(02)00006-3
9. Martinez-Alcaraz, P., de la Barra, P., Andriotis, C.P., et al.: Personalised building controls based on individual thermal preferences for energy efficiency and thermal comfort. In: Berardi, U. (ed.) Multiphysics and Multiscale Building Physics, pp. 203–210. Springer, Singapore (2025)
10. Parkinson, T., et al.: ASHRAE global database of thermal comfort field measurements. Dryad (2022)
11. Gelman, A., Hwang, J., Vehtari, A.: Understanding predictive information criteria for Bayesian models (2013)