NATIONAL TECHNICAL UNIVERSITY OF ATHENS School of Electrical Engineering and Computer Science Division of Communications, Electronics and Information Systems



Modeling and Personalization of Thermal Comfort Using Machine Learning and Transfer Learning Approaches

Odysseas Moschidis

Supervisor: Emmanouel Varvarigos Professor, ECE NTUA

A thesis submitted to the National Technical University of Athens in partial fulfillment of the requirements for the degree of electrical engineering and computer science

Athens, June 2025

NATIONAL TECHNICAL UNIVERSITY OF ATHENS School of Electrical Engineering and Computer Science Division of Communications, Electronics and Information Systems



Modeling and Personalization of Thermal Comfort Using Machine Learning and Transfer Learning Approaches

Odysseas Moschidis

Supervisor: Emmanouel Varvarigos Professor, ECE NTUA

Approved by the three-member scientific committee on June 17, 2025 $\,$

Emmanouil Varvarigos Professor, ECE NTUA

Theodora Varvarigou Professor, ECE NTUA

Hercules Avramopoulos Professor, ECE NTUA

Athens, June 2025

Odysseas Moschidis

Graduate of School of Electrical and Computer Engineering, National Technical University of Athens

Copyright © Odysseas Moschidis, 2025 All rights reserved.

You may not copy, reproduce, distribute, publish, display, modify, create derivative works, transmit, or in any way exploit this thesis or part of it for commercial purposes. You may reproduce, store or distribute this thesis for non-profit educational or research purposes, provided that the source is cited, and the present copyright notice is retained. Inquiries for commercial use should be addressed to the original author.

The ideas and conclusions presented in this paper are the author's and do not necessarily reflect the official views of the National Technical University of Athens.

Abstract

Thermal comfort is a multifactorial phenomenon influenced by environmental, physiological, and contextual factors. As smart building environments evolve toward personalized control systems, data-driven models have emerged as key enablers of adaptive comfort prediction. However, existing models often struggle to generalize across users, time, and deployment contexts.

This thesis investigates the problem of user preference transferability and dynamic comfort modeling using real-world datasets enriched with multivariate sensor data. It emphasizes user-centric modeling through feature engineering, temporal embeddings, and the use of sequential deep learning architectures. A particular focus is placed on exploring CNN-LSTM networks, which can jointly learn spatial representations and temporal dependencies inherent in thermal comfort signals.

Additionally, the study explores model adaptation techniques—including instance reweighting and feature distribution alignment—that aim to improve generalization across users and climates. These approaches are evaluated with respect to their capacity to transfer learned comfort representations while minimizing degradation in predictive performance.

In evaluating the predictive models, emphasis is placed on real-world deployment scenarios, including model simplicity, data sparsity, and on-device computation limits. These considerations help ensure that the insights derived from this research remain applicable to practical smart building implementations and personalized control systems.

Keywords: Thermal Comfort, Personalized Modeling, Smart Environments, Machine Learning, Transfer Learning, User Preferences, Sensor Data Fusion, Random Forest, Exploratory Data Analysis, Human-Centered AI

Περίληψη

Η θερμική άνεση αποτελεί ένα πολυπαραγοντικό φαινόμενο που επηρεάζεται από περιβαλλοντικούς, φυσιολογικούς και συγκειμενικούς παράγοντες. Καθώς τα smart buildings εξελίσσονται προς εξατομικευμένα συστήματα ελέγχου, τα υποδείγματα βασισμένα σε δεδομένα αναδεικνύονται ως βασικά εργαλεία για την προσαρμοστική πρόβλεψη της άνεσης. Ωστόσο, τα υπάρχοντα υποδείγματα δυσκολεύονται συχνά να γενικεύσουν μεταξύ χρηστών, χρονικών περιόδων και πλαισίων υλοποίησης.

Η παρούσα εργασία διερευνά το ζήτημα της μεταφερσιμότητας των προτιμήσεων των χρηστών και της δυναμικής μοντελοποίησης της θερμικής άνεσης, αξιοποιώντας δεδομένα από πραγματικά περιβάλλοντα εμπλουτισμένα με πολυμεταβλητά αισθητήρια σήματα. Δίνεται έμφαση στη μοντελοποίηση με επίκεντρο τον χρήστη μέσω της μηχανικής χαρακτηριστικών, της ενσωμάτωσης χρονικής πληροφορίας και της χρήσης ακολουθιακών αρχιτεκτονικών βαθιάς μάθησης. Ιδιαίτερη έμφαση δίνεται στη χρήση CNN-LSTM δικτύων, τα οποία μπορούν να μάθουν από κοινού χωρικές αναπαραστάσεις και χρονικές εξαρτήσεις που χαρακτηρίζουν τα σήματα θερμικής άνεσης.

Επιπλέον, η μελέτη εξετάζει τεχνικές προσαρμογής μοντέλων—όπως αναπροσαρμογή βαρών δειγμάτων και ευθυγράμμιση κατανομών χαρακτηριστικών—με στόχο τη βελτίωση της γενίκευσης μεταξύ χρηστών και κλιματικών συνθηκών. Οι προσεγγίσεις αυτές αξιολογούνται ως προς την ικανότητά τους να μεταφέρουν εκπαιδευμένες αναπαραστάσεις άνεσης χωρίς σημαντική απώλεια ακρίβειας πρόβλεψης.

Κατά την αξιολόγηση των προβλεπτικών μοντέλων, δίνεται ιδιαίτερη έμφαση σε σενάρια πραγματικής εφαρμογής, λαμβάνοντας υπόψη την απλότητα των μοντέλων, τη σπανιότητα δεδομένων και τους περιορισμούς επεξεργασίας επί συσκευής. Αυτές οι παραδοχές διασφαλίζουν ότι τα συμπεράσματα της παρούσας έρευνας μπορούν να εφαρμοστούν σε ρεαλιστικά smart building systems και συστήματα προσωποποιημένου ελέγχου.

Λέξεις-κλειδιά: Θερμική Άνεση, Εξατομικευμένη Μοντελοποίηση, Έξυπνα Περιβάλλοντα, Μηχανική Μάθηση, Μεταφορά Μάθησης, Προτιμήσεις Χρηστών, Ενοποίηση Αισθητήριων Δεδομένων, Τυχαίο Δάσος, Εξερευνητική Ανάλυση Δεδομένων, Ανθρωποκεντρική Τεχνητή Νοημοσύνη

Acknowledgments

I would like to thank my advisor, Professor Emmanouil Varvarigos, for his guidance and support during the preparation of this thesis. His suggestions and comments were valuable in helping me stay focused and develop the core ideas of this work. Also I would like to thank my friend and supervisor Michael for his support and camaraderie.

I would also like to extend heartfelt thanks to my family and friends, whose encouragement and understanding sustained me throughout this journey. A special thank you to my mother, whose patience, care, and quiet strength provided unwavering support during the most demanding moments. And my sister Kate Mos, for her unwavering belief in me and wonderful hospitality.

Contents

\mathbf{A}	bstra	let	5
П	ερίλη	ηψη	7
Li	st of	Tables	14
Li	st of	Figures	15
E:	 Εισο Αναι Θεω Πειρ Συζι Μελ Intr 1.2 1.3 1.4 	ής Περίληψη στα Ελληνικά αγωγή	 16 16 17 18 19 20 21 23 24 25 26 27
2	 1.3 Lite 2.1 2.2 2.3 	Prature Review Foundations of Thermal Comfort Modeling 2.1.1 Physical Models: PMV and PPD 2.1.2 Adaptive Comfort Models 2.1.3 Psychological and Behavioral Extensions 2.1.4 Personalization and User-Centric Modeling 2.2.1 Personalized Comfort Models (PCMs) 2.2.2 Challenges in Generalization 2.2.3 Longitudinal Studies and Feedback Systems Machine Learning Approaches in Comfort Modeling 2.3.1 Supervised Learning Techniques 2.3.2 Unsupervised Learning and Clustering 2.3.3 Temporal Models and Deep Architectures	29 29 30 30 31 31 31 32 32 32 33 33

		2.3.4	Reinforcement and Active Learning				
	2.4	Transf	fer Learning and Domain Adaptation				
		2.4.1	Domain Transfer Challenges				
		2.4.2	Adaptation Techniques				
		2.4.3	Hybrid Models				
	2.5	Integr	ation into Smart Systems				
		2.5.1	IoT and Wearables				
		2.5.2	Edge Deployment and Privacy				
		2.5.3	Interface Design and User Interaction				
3	The	oretic	al Background 37				
Č	3.1	Found	ations of Thermal Comfort Modeling 37				
	0.1	3.1.1	Predicted Mean Vote (PMV) and Predicted Percentage of Dissat-				
		0.1.1	isfied (PPD) 37				
		312	Adaptive Comfort Model 38				
		313	Data-Driven Comfort Models 39				
		3.1.0	Cognitive and Psychological Models of Comfort 40				
	3.2	Machi	ne Learning Fundamentals 40				
	0.2	321	Introduction to Machine Learning 40				
		3.2.2	Principles of Machine Learning 40				
		323	Types of Machine Learning 41				
		3.2.4	Training Machine Learning Models 41				
		3.2.5	Applications of Machine Learning				
		3.2.6	Challenges and Future Directions				
	3.3	Machi	ne Learning and Data-Driven Techniques				
	0.0	3.3.1	Supervised Learning 43				
		3.3.2	Unsupervised Learning and Dimensionality Reduction				
		3.3.3	Reinforcement Learning				
		3.3.4	Transfer Learning				
	3.4	Featur	re Representation and Engineering				
		3.4.1	User-Centric Data				
		3.4.2	Environmental Features				
		3.4.3	Contextual and Interaction Features				
		3.4.4	Time Resolution and Historical Modeling				
4	Exr	erime	nts and Results 55				
-	4 1	Datas	et Overview 55				
	4.2 Exploratory Data Analysis						
	1.2	4.2.1	Target Variable: Thermal Sensation 56				
		4.2.2	Feature Correlation Analysis				
		4.2.3	Principal Component Analysis (PCA) 60				
		4.2.4	User-Level Record Distribution 62				
		4.2.5	Per-User Correlation Distribution Analysis				
		4.2.6	Feature Engineering				
	4.3	Baseli	ne Models				

	4.4	Experimental Results	68
		4.4.1 Model Performance Overview	68
		4.4.2 Logistic Regression	69
		4.4.3 Random Forest	69
		4.4.4 Multi-Layer Perceptron (MLP)	70
		4.4.5 XGBoost	70
		4.4.6 Confusion Matrix Analysis	71
		4.4.7 Effect of Window Opening on Predicted Comfort	74
		4.4.8 Geographical Subsetting and Model Performance	78
	4.5	Transfer Learning for Climate-Adaptive	
		Modeling	80
		4.5.1 Model Design and Implementation	80
		4.5.2 Training and Evaluation on the Source Domain	81
		4.5.3 Transfer Learning to Target Climate	83
		4.5.4 Discussion and Implications	86
		•	
5	Dis	cussion	88
	5.1	Summary of Modeling Results	88
	5.2	Limitations of Model Performance	89
	5.3	Practical Implications for Smart Environments	91
	5.4	Personalization Potential and User-Centric	
		Modeling	92
G	E.t.	une Wenk	04
0	ги. 61	Enhancing Comfort Prediction through Dichor Data Sources	94 04
	0.1 6.2	Transfor Learning and Domain Adaptation	94 05
	0.2	6.2.1 Motivation for Transfer Learning	95 05
		6.2.2 Domain Adaptation Techniques	95 05
		6.2.2 Domain Adaptation Techniques	95 05
		6.2.4 Eutrop Directions	90 06
	63	Clustering and Segmentation for Descendized	90
	0.0	Modoling	96
		6.3.1 Importance of User Segmentation	90 06
		6.3.2 Clustoring Tochniques	90 06
		6.3.2 Implementation in Smart Environments	90 07
		6.3.4 Challenges and Considerations	97 07
	64	Model Simplification and Edge Deployment	91
	0.4 6 5	Hybrid Physical Statistical Modeling	00
	0.0 6 6	Active Learning and Feedback-Efficient Personalization	99 100
	0.0 6 7	Real World System Integration	100
	0.7		101
7	Ref	erences	103

List of Tables

4.1	Performance Summary of Baseline Models		68
-----	--	--	----

List of Figures

4.1	Correlation matrix of numerical features	58
4.2	Correlation of selected features with thermal sensation	59
4.3	Cumulative variance explained by PCA components	61
4.4	Number of users with at least N records	62
4.5	Distributions of per-user feature correlations with thermal sensation	64
4.6	Cross-Model Comparison Bar Plot	69
4.7	Random Forest Feature Importance	70
4.8	Confusion Matrix – Logistic Regression	71
4.9	Confusion Matrix – Random Forest	72
4.10	Confusion Matrix – MLP	73
4.11	Confusion Matrix – XGBoost	74
4.12	Feature importance from Random Forest model including the window vari-	
	able	76
4.13	Change in predicted comfort (Δ Comfort) under window state alteration	
	$(counterfactual test) \dots \dots \dots \dots \dots \dots \dots \dots \dots $	76
4.14	Change in distance to neutral comfort score $(Comfort_{closed} - Comfort_{open})$.	
	Positive values indicate improved proximity to thermal neutrality when	
	window is open	77
4.15	Training vs. Validation Loss across epochs for CNN-LSTM on ASHRAE	
	dataset	83
4.16	Training vs. Validation Accuracy across epochs for CNN-LSTM on ASHRAE	
	dataset	84
4.17	Training vs. Validation Loss – Transfer Learning on Tropical Savanna	
	Climate	85
4.18	Training vs. Validation Accuracy – Transfer Learning on Tropical Savanna	
	Climate	86

Εκτενής Περίληψη στα Ελληνικά

Εισαγωγή

Η θερμική άνεση αποτελεί ένα πολυπαραγοντικό και υποκειμενικό φαινόμενο, το οποίο εξαρτάται από ένα πλήθος περιβαλλοντικών, φυσιολογικών και συμπεριφορικών παραμέτρων. Η κατανόηση και η ακριβής πρόβλεψη της θερμικής άνεσης των χρηστών είναι κρίσιμη για τον σχεδιασμό ευφυών κτιριακών περιβαλλόντων που στοχεύουν τόσο στην ευημερία των ενοίκων όσο και στη βελτιστοποίηση της ενεργειακής κατανάλωσης.

Εξέλιξη των μοντέλων θερμικής άνεσης

Παραδοσιαχά μοντέλα, όπως το PMV/PPD, βασίζονται σε φυσιολογιχά χαι θερμοδυναμιχά μοντέλα του ανθρώπινου οργανισμού. Παρόλο που προσφέρουν ένα θεωρητιχό πλαίσιο, παρουσιάζουν σημαντιχούς περιορισμούς σε δυναμιχά ή εξατομιχευμένα περιβάλλοντα. Οι προσεγγίσεις αυτές δεν λαμβάνουν υπόψη τις υποχειμενιχές προτιμήσεις ή την προσαρμογή του χρήστη, γεγονός που οδηγεί σε συχνές αποχλίσεις μεταξύ προβλεπόμενης χαι πραγματιχής θερμιχής αίσθησης.

Η σημασία της εξατομίχευσης στα έξυπνα περιβάλλοντα

Η μετάβαση προς ευφυή κτίρια και περιβάλλοντα που ενσωματώνουν τεχνολογίες ΙοΤ και φθαρτούς αισθητήρες επιτρέπει τη συνεχή παρακολούθηση δεδομένων από πολλαπλές πηγές — από περιβαλλοντικά στοιχεία έως φυσιολογικά σήματα. Αυτό προσφέρει μια μοναδική ευκαιρία για εξατομίκευση, καθιστώντας δυνατή τη δημιουργία μοντέλων που ανταποκρίνονται στις ιδιαίτερες ανάγκες κάθε χρήστη.

Πρόκληση της μεταφοράς προτιμήσεων

Η βασική πρόκληση που διερευνάται σε αυτή τη διατριβή αφορά την ικανότητα των μοντέλων θερμικής άνεσης να προσαρμόζονται και να μεταφέρονται μεταξύ διαφορετικών χρηστών και περιβαλλόντων. Οι ατομικές προτιμήσεις μεταβάλλονται δυναμικά και εξαρτώνται από το πλαίσιο, γεγονός που καθιστά δύσκολη τη γενίκευση. Αυτό οδηγεί στην ανάγκη για τεχνικές transfer learning, προσαρμογής παραμέτρων και δυναμικής ανατροφοδότησης.

Στόχοι και πεδίο της διατριβής

Η παρούσα εργασία επικεντρώνεται στη μελέτη και αξιολόγηση τεχνικών μάθησης, όπως οι CNN-LSTM αρχιτεκτονικές για σπατιοχρονική μάθηση, και τεχνικές μεταφοράς και ευθυγράμμισης χαρακτηριστικών, με σκοπό τη βελτίωση της εξατομικευμένης πρόβλεψης θερμικής άνεσης. Η έρευνα βασίζεται σε αναλύσεις υπαρχόντων δεδομένων και αξιολογεί τη δυνατότητα απλοποίησης μοντέλων για υλοποίηση σε συσκευές αιχμής.

Περιορισμοί της μελέτης

Το πεδίο της διατριβής περιορίζεται σε ανάλυση υπαρχόντων δεδομένων από τη ASHRAE Global Thermal Comfort Database II και δεν περιλαμβάνει την ανάπτυξη φυσικού εξοπλισμού ή ευρείας κλίμακας πειραματικές υλοποιήσεις. Ωστόσο, τα αποτελέσματα της μελέτης προσφέρουν σημαντικές ενδείξεις για την κατεύθυνση μελλοντικής έρευνας και ανάπτυξης σε έξυπνα περιβάλλοντα.

Με αυτό το υπόβαθρο, η εργασία επιδιώχει να συμβάλλει τόσο σε θεωρητικό όσο χαι σε εφαρμοσμένο επίπεδο, με στόχο την ανάπτυξη προσαρμοστικών, εξατομικευμένων και τεχνολογικά υλοποιήσιμων μοντέλων θερμικής άνεσης.

Ανασκόπηση Βιβλιογραφίας

Το παρόν κεφάλαιο παρέχει μια συστηματική επισκόπηση των θεωρητικών και τεχνολογικών προσεγγίσεων που σχετίζονται με τη μοντελοποίηση της θερμικής άνεσης σε έξυπνα περιβάλλοντα. Εξετάζονται οι παραδοσιακές και σύγχρονες μέθοδοι, οι τεχνικές μηχανικής μάθησης, οι στρατηγικές εξατομίκευσης και οι προσεγγίσεις ενσωμάτωσης σε συστήματα ΙοΤ.

Μοντέλα Θερμικής Άνεσης

Η ενότητα ξεκινά με αναφορά στα φυσικά μοντέλα PMV/PPD, τα οποία βασίζονται σε φυσιολογικές παραμέτρους του ανθρώπινου οργανισμού (Fanger 1970). Εξηγούνται τα μειονεκτήματα των μοντέλων αυτών σε δυναμικά ή μη ελεγχόμενα περιβάλλοντα. Ακολουθεί η παρουσίαση των προσαρμοστικών μοντέλων (Adaptive Comfort Models), τα οποία λαμβάνουν υπόψη τη συμπεριφορική προσαρμογή και την επιρροή του περιβάλλοντος. Τέλος, αναλύονται οι ψυχολογικές και συμπεριφορικές επεκτάσεις, οι οποίες αποτελούν γέφυρα προς τα προσωποποιημένα μοντέλα.

Εξατομίκευση και Μοντελοποίηση Προτιμήσεων

Η βιβλιογραφία σε εξατομιχευμένα μοντέλα θερμιχής άνεσης (Personalized Comfort Models) εξετάζει χαραχτηριστιχά όπως το φύλο, η ηλιχία, η BMI, χαι φυσιολογιχά σήματα. Ειδιχή αναφορά γίνεται σε μελέτες που χρησιμοποίησαν δεδομένα από φθαρτούς αισθητήρες (Lee and Chun 2021, Jayathissa et al. 2020). Επισημαίνονται οι δυσχολίες γενίχευσης χαι η μεγάλη ετερογένεια προτιμήσεων μεταξύ χρηστών. Επιπλέον, παρουσιάζονται μαχροχρόνιες μελέτες με συστήματα ανατροφοδότησης που προσαρμόζονται δυναμιχά με την πάροδο του χρόνου (Tekler et al. 2023, Gnecco et al. 2023).

Τεχνικές Μηχανικής Μάθησης

Η ενότητα αυτή εστιάζει στις εφαρμογές επιβλεπόμενης μάθησης όπως οι λογιστικές παλινδρομήσεις, τα δέντρα απόφασης, τα MLPs, και οι συναθροιστικές μέθοδοι (ensemble models). Επιπλέον, εξετάζονται αλγόριθμοι μη επιβλεπόμενης μάθησης και τεχνικές συμπίεσης όπως PCA, k-means και DBSCAN, οι οποίες βοηθούν στον εντοπισμό λανθανουσών ομάδων προτιμήσεων. Αναλύονται επίσης χρονικές αρχιτεκτονικές όπως τα CNN-LSTM και εν συντομία οι προσεγγίσεις ενισχυτικής και ενεργής μάθησης.

Μεταφορά Μάθησης και Προσαρμογή Πεδίων

Η δυνατότητα γενίκευσης μοντέλων μεταξύ διαφορετικών κλιμάτων ή πληθυσμών παραμένει περιορισμένη, όπως δείχνουν οι μελέτες των Yang et al. (2025) και Gao et al. (2021). Εξετάζονται τεχνικές προσαρμογής όπως η εκ νέου σταθμολόγηση παραδειγμάτων (instance reweighting), η ευθυγράμμιση χαρακτηριστικών (MMD) και η εχθρική εκπαίδευση (adversarial training). Τέλος, αναφέρονται τα υβριδικά μοντέλα που συνδυάζουν εξόδους φυσικών μοντέλων με στατιστική επεξεργασία (Zhou et al. 2021).

Ενσωμάτωση σε Έξυπνα Συστήματα

Η πρόοδος στην τεχνολογία IoT και των φθαρτών αισθητήρων επιτρέπει τη συνεχή συλλογή δεδομένων (Liu 2018). Γίνεται αναφορά στη σημασία των WELL Building Standards για τη διασφάλιση ποιότητας και άνεσης. Επιπλέον, εξετάζεται η ανάγκη για ελαφριά μοντέλα που μπορούν να υλοποιηθούν σε συσκευές αιχμής (edge devices) μέσω τεχνικών όπως η περικοπή (pruning) και η ποσοτικοποίηση (quantization) (Liang 2021, Francy 2024). Τέλος, αναλύεται ο ρόλος του σχεδιασμού διεπαφής και της ανθρώπινης αλληλεπίδρασης για τη χρηστικότητα των προσωποποιημένων συστημάτων (Zhu et al. 2021).

Η ανασκόπηση αυτή τοποθετεί την παρούσα εργασία στο επιστημονικό και τεχνολογικό της πλαίσιο, αναδεικνύοντας τα ερευνητικά κενά και τις ευκαιρίες για προχωρημένες και υλοποιήσιμες προσεγγίσεις μοντελοποίησης της θερμικής άνεσης.

Θεωρητικό Υπόβαθρο

Το παρόν κεφάλαιο παρουσιάζει το θεωρητικό υπόβαθρο που υποστηρίζει τη μελέτη της μεταφερσιμότητας και ανάλυσης προτιμήσεων χρηστών σε πραγματικά έξυπνα περιβάλλοντα. Εστιάζει σε βασικές έννοιες και τεχνολογίες που σχετίζονται με την εξατομικευμένη θερμική άνεση, την προσαρμογή συστημάτων σε προτιμήσεις χρηστών και τη διαχείριση δεδομένων σε έξυπνα περιβάλλοντα.

Θερμική Άνεση και Εξατομίκευση

Η θερμική άνεση είναι μια υποκειμενική εμπειρία που επηρεάζεται από περιβαλλοντικούς παράγοντες (όπως θερμοκρασία, υγρασία, ροή αέρα) και ατομικά χαρακτηριστικά (όπως μεταβολισμός, ένδυση). Παραδοσιακά μοντέλα, όπως το Predicted Mean Vote (PMV), παρέχουν γενικευμένες εκτιμήσεις αλλά δεν λαμβάνουν υπόψη την ατομική διαφοροποίηση. Η ανάγκη για εξατομικευμένα μοντέλα έχει οδηγήσει στην ανάπτυξη προσεγγίσεων που ενσωματώνουν δεδομένα από αισθητήρες και προτιμήσεις χρηστών για την παροχή προσωποιημένων εμπειριών θερμικής άνεσης.

Έξυπνα Περιβάλλοντα και Τεχνολογίες ΙοΤ

Τα έξυπνα περιβάλλοντα (smart environments) ενσωματώνουν τεχνολογίες Internet of Things (IoT) για τη συλλογή και ανάλυση δεδομένων σε πραγματικό χρόνο. Οι αισθητήρες καταγράφουν πληροφορίες όπως θερμοκρασία, υγρασία, παρουσία χρηστών, ενώ οι ενεργοποιητές (actuators) επιτρέπουν την αυτόματη προσαρμογή του περιβάλλοντος. Η διαχείριση αυτών των δεδομένων απαιτεί προηγμένες τεχνικές επεξεργασίας και ανάλυσης για την εξαγωγή χρήσιμων γνώσεων και την υποστήριξη αποφάσεων.

Μηχανική Μάθηση και Μοντελοποίηση Προτιμήσεων

Η μηχανική μάθηση παρέχει εργαλεία για την ανάλυση μεγάλων ποσοτήτων δεδομένων και την αναγνώριση προτύπων. Στο πλαίσιο της θερμικής άνεσης, χρησιμοποιούνται τεχνικές όπως τα Convolutional Neural Networks (CNN) για την εξαγωγή χωρικών χαρακτηριστικών και τα Long Short-Term Memory (LSTM) δίκτυα για την ανάλυση χρονικών εξαρτήσεων. Οι συνδυασμένες αρχιτεκτονικές CNN-LSTM επιτρέπουν την ταυτόχρονη μάθηση χωρικών και χρονικών προτύπων, βελτιώνοντας την ακρίβεια των προβλέψεων θερμικής άνεσης.

Μεταφορά Μάθησης και Προσαρμογή Μοντέλων

Η μεταφορά μάθησης (transfer learning) επιτρέπει τη χρήση γνώσης από ένα περιβάλλον ή χρήστη για τη βελτίωση της απόδοσης σε ένα άλλο. Τεχνικές όπως η ευθυγράμμιση χαρακτηριστικών (feature alignment) και η επαναστάθμιση παραδειγμάτων (instance reweighting) βοηθούν στην προσαρμογή μοντέλων σε νέους χρήστες ή περιβάλλοντα με περιορισμένα δεδομένα. Αυτές οι προσεγγίσεις είναι κρίσιμες για την ανάπτυξη ευέλικτων και προσαρμοστικών συστημάτων θερμικής άνεσης. Και αυτό για τον Μπίλι Νιάμου.

Προστασία Προσωπικών Δεδομένων και Διαχείριση Ταυτότητας

Η συλλογή και ανάλυση προσωπικών δεδομένων σε έξυπνα περιβάλλοντα εγείρει ζητήματα ιδιωτικότητας και ασφάλειας. Τεχνολογίες όπως η ενισχυμένη διαχείριση ταυτότητας (Privacy-Enhancing Identity Management, PE-IdM) προσφέρουν μηχανισμούς για τον έλεγχο της ροής προσωπικών πληροφοριών και την προστασία της ιδιωτικότητας των χρηστών. Η ενσωμάτωση αυτών των τεχνολογιών είναι απαραίτητη για την αποδοχή και εμπιστοσύνη των χρηστών σε έξυπνα συστήματα.

Η κατανόηση αυτών των θεωρητικών εννοιών και τεχνολογιών είναι απαραίτητη για την ανάπτυξη συστημάτων που μπορούν να προσαρμόζονται στις ατομικές προτιμήσεις χρηστών, να λειτουργούν αποτελεσματικά σε διαφορετικά περιβάλλοντα και να διασφαλίζουν την προστασία των προσωπικών δεδομένων.

Πειράματα και Αποτελέσματα

Το παρόν χεφάλαιο παρουσιάζει τη μεθοδολογία και τα κύρια αποτελέσματα των πειραμάτων που διεξήχθησαν για την αξιολόγηση της μεταφερσιμότητας και της ανάλυσης προτιμήσεων χρηστών σε έξυπνα περιβάλλοντα. Η εστίαση βρίσκεται στην εφαρμογή τεχνικών μηχανικής μάθησης για την πρόβλεψη της θερμικής άνεσης, καθώς και στη διερεύνηση της ικανότητας μοντέλων να προσαρμόζονται σε διαφορετικούς χρήστες και συνθήκες.

Πειραματική Διαδικασία

Η ανάλυση βασίστηκε σε προϋπάρχοντα δεδομένα που συλλέχθηκαν από αισθητήρες περιβάλλοντος και δεδομένα χρήστη. Πραγματοποιήθηκε διερεύνηση χαρακτηριστικών μέσω στατιστικής ανάλυσης και τεχνικών όπως η Principal Component Analysis (PCA). Στη συνέχεια, εφαρμόστηκαν βασικά εποπτευόμενα μοντέλα μηχανικής μάθησης όπως Logistic Regression, Random Forest, Multi-Layer Perceptron (MLP) και XGBoost για την πρόβλεψη του θερμικού αισθήματος.

Αποτελέσματα και Παρατηρήσεις

Τα αποτελέσματα έδειξαν ότι τα απλά μοντέλα μπορούν να αποδώσουν ικανοποιητικά υπό ορισμένες συνθήκες, ειδικά όταν πραγματοποιείται κατάλληλη επεξεργασία χαρακτηριστικών. Εντοπίστηκε διαφοροποίηση στις επιδόσεις ανά χρήστη και περιβάλλον, γεγονός που ενισχύει την ανάγκη για εξατομίκευση. Η εφαρμογή μεταφοράς μάθησης δοκιμάστηκε σε περιορισμένο βαθμό, με ενδείξεις ότι η χρήση τεχνικών όπως η ευθυγράμμιση χαρακτηριστικών και η επαναστάθμιση μπορεί να προσφέρει βελτιώσεις, ιδιαίτερα όταν υπάρχει διαφορά κλίματος ή χρήστη μεταξύ των πηγών δεδομένων.

Συνοπτικά Συμπεράσματα

Τα ευρήματα της πειραματικής μελέτης προσφέρουν ενδείξεις για τις δυνατότητες εφαρμογής εξατομικευμένων μοντέλων σε έξυπνα περιβάλλοντα, αν και η επίδοση επηρεάζεται σημαντικά από την ποιότητα και ποικιλομορφία των δεδομένων. Η αξιολόγηση της μεταφερσιμότητας υποδεικνύει την ανάγκη για περαιτέρω μελέτη και βελτίωση των μεθόδων προσαρμογής, ώστε να επιτυγχάνεται καλύτερη γενίκευση μεταξύ διαφορετικών καταστάσεων.

Γενικά, το κεφάλαιο αυτό σκιαγραφεί ένα ρεαλιστικό πλαίσιο αξιολόγησης, παρέχοντας τις βάσεις για πιο προχωρημένες μεθόδους μοντελοποίησης που θα μπορούσαν να ενσωματωθούν σε πλήρως λειτουργικά, χρήστη-κεντρικά συστήματα θερμικής άνεσης στο μέλλον.

Συζήτηση

Το κεφάλαιο αυτό επιχειρεί μια συνολική αποτίμηση των ευρημάτων που προέκυψαν από την ανάλυση και τα πειράματα του προηγούμενου κεφαλαίου. Εξετάζονται τα πλεονεκτήματα, οι περιορισμοί και οι πρακτικές επιπτώσεις των μεθόδων που εφαρμόστηκαν, με στόχο την καλύτερη κατανόηση της δυνατότητας εξατομίκευσης και προσαρμογής μοντέλων θερμικής άνεσης σε πραγματικά έξυπνα περιβάλλοντα.

Αποτίμηση Μοντέλων και Μεθόδων

Οι εφαρμοσμένες τεχνικές μηχανικής μάθησης ανέδειξαν ότι ακόμη και απλά μοντέλα μπορούν να επιτύχουν αποδεκτά αποτελέσματα υπό συγκεκριμένες συνθήκες. Η διερεύνηση ανά χρήστη και γεωγραφική υποομάδα ανέδειξε σαφείς αποκλίσεις στις θερμικές προτιμήσεις και επιβεβαίωσε την ανάγκη για εξατομικευμένη προσέγγιση. Οι τεχνικές μεταφοράς μάθησης, αν και περιορισμένες σε εφαρμογή, παρείχαν υποσχόμενα αποτελέσματα όσον αφορά την προσαρμοστικότητα.

Περιορισμοί και Επιφυλάξεις

Ένας από τους χύριους περιορισμούς αφορά την ποιότητα χαι την ετερογένεια των δεδομένων. Η απουσία πραγματιχού χρόνου παρεμβάσεων χαι οι περιορισμοί στη συλλογή φυσιολογιχών δεδομένων περιορίζουν το εύρος της εφαρμοσιμότητας των μοντέλων. Επίσης, η έλλειψη εκτεταμένων πραγματιχών δοχιμών μειώνει την αξιοπιστία των συμπερασμάτων σε επιχειρησιαχό περιβάλλον.

Πρακτικές Επιπτώσεις και Εφαρμογές

Τα αποτελέσματα δείχνουν ότι τα έξυπνα περιβάλλοντα μπορούν να επωφεληθούν σημαντικά από μοντέλα που ενσωματώνουν δυναμική και εξατομικευμένη πρόβλεψη θερμικής άνεσης. Ιδιαίτερα, η δυνατότητα απλοποίησης των μοντέλων για υλοποίηση σε edge συσκευές μπορεί να προσφέρει πρακτικά οφέλη σε ό,τι αφορά την ταχύτητα απόκρισης, την ιδιωτικότητα και την ενεργειακή απόδοση.

Συμπεράσματα

Η ανάλυση των ευρημάτων καταδεικνύει ότι η κατεύθυνση της εξατομίκευσης στη θερμική άνεση είναι ουσιαστική και πολλά υποσχόμενη. Ωστόσο, για να επιτευχθεί πλήρης αξιοποίηση της δυναμικής αυτής απαιτούνται πιο εντατικές μελέτες, ποιοτικότερα δεδομένα και πραγματικές δοκιμές σε έξυπνα κτίρια.

Μελλοντική Έρευνα

Το τελευταίο κεφάλαιο της εργασίας επικεντρώνεται σε προτάσεις για μελλοντική έρευνα με στόχο την αντιμετώπιση των περιορισμών που εντοπίστηκαν και τη διεύρυνση της αποτελεσματικότητας των μοντέλων θερμικής άνεσης σε πραγματικά έξυπνα περιβάλλοντα. Οι ερευνητικές κατευθύνσεις οργανώνονται σε θεματικές ενότητες που σχετίζονται με την ενίσχυση της απόδοσης, την εξατομίκευση, την υλοποίηση σε πραγματικά συστήματα και τη μείωση της εμπλοκής του χρήστη.

Ενίσχυση της Πρόβλεψης μέσω Πλουσιότερων Δεδομένων

Η ενσωμάτωση φυσιολογικών δεδομένων, όπως η θερμοκρασία δέρματος και η μεταβλητότητα καρδιακού ρυθμού, μπορεί να ενισχύσει σημαντικά την εξατομίκευση των μοντέλων. Επιπλέον, παράγοντες όπως οι συγκεντρώσεις Ό₂ και η συμπεριφορά του χρήστη (π.χ. άνοιγμα παραθύρων) συνιστούν κρίσιμα περιβαλλοντικά και συμφραζόμενα χαρακτηριστικά.

Μεταφορά Μάθησης και Προσαρμογή Τομέων

Η μελλοντική εργασία μπορεί να εξερευνήσει πιο σύνθετες τεχνικές προσαρμογής, όπως η αντιπαραθετική μάθηση και οι μέθοδοι βασισμένες στη μέγιστη απόκλιση μέσου όρου. Αυτές επιτρέπουν την καλύτερη προσαρμογή μοντέλων μεταξύ χρηστών ή κλιματικών περιοχών, χωρίς σημαντική απώλεια απόδοσης.

Ομαδοποίηση και Τμηματοποίηση

Η τμηματοποίηση των χρηστών σε ομάδες με παρόμοιες προτιμήσεις μπορεί να βελτιώσει την αποδοτικότητα και την ευρωστία των μοντέλων, ιδίως όταν τα δεδομένα είναι περιορισμένα. Η εφαρμογή τεχνικών όπως οι συστάδες (clustering) αναδεικνύεται ως ιδιαίτερα σημαντική.

Απλοποίηση Μοντέλων και Υλοποίηση σε Edge Συσκευές

Η διερεύνηση τεχνικών συμπίεσης, όπως η κλάδευση και η ποσοτικοποίηση, ενισχύει τη δυνατότητα ενσωμάτωσης μοντέλων σε συσκευές χαμηλής ισχύος όπως έξυπνοι θερμοστάτες και ελεγκτές Η[°]Α[°]. Αυτή η υλοποίηση μπορεί να προσφέρει πραγματικό χρόνο αποκρίσεις χωρίς εξάρτηση από το cloud.

Υβριδικά Μοντέλα Φυσικής και Στατιστικής

Ο συνδυασμός παραδοσιακών φυσικών μοντέλων, όπως το ΠΜ[°], με στατιστικά μοντέλα επιτρέπει τη σύζευξη θεωρητικής αυστηρότητας και εμπειρικής προσαρμοστικότητας. Τέτοιες προσεγγίσεις μπορούν να οδηγήσουν σε πιο ερμηνεύσιμα και ταυτόχρονα ακριβή μοντέλα.

Μάθηση με Ενεργό Χρήστη και Μείωση Ανατροφοδότησης

Οι τεχνικές ενεργής μάθησης μπορούν να μειώσουν τον φόρτο εισόδου από τον χρήστη, ζητώντας σχόλια μόνο όταν αυτά είναι πληροφοριακά. Επιπλέον, η αξιοποίηση φυσιολογικών ενδείξεων επιτρέπει παθητική συλλογή ανατροφοδότησης.

Ενσωμάτωση σε Πραγματικά Συστήματα

Η μελλοντική έρευνα θα πρέπει να επικεντρωθεί στην υλοποίηση και δοκιμή των συστημάτων σε πραγματικά περιβάλλοντα, αξιολογώντας τη χρηστικότητα, την ικανοποίηση των χρηστών και τη μακροπρόθεσμη ενεργειακή επίδραση. Παράλληλα, η διασφάλιση της ιδιωτικότητας και της διαλειτουργικότητας παραμένει κρίσιμη πρόκληση.

Chapter 1 Introduction

1.1 Background and Motivation

Thermal comfort is a multifactorial phenomenon shaped by a complex interplay of environmental, physiological, behavioral, and psychological variables. Classical definitions often emphasize its subjective nature—commonly understood as "that condition of mind which expresses satisfaction with the thermal environment," as defined by ASHRAE. To model and assess thermal comfort, researchers have historically relied on physical models that aggregate environmental metrics such as air temperature, humidity, air velocity, and mean radiant temperature, alongside metabolic rate and clothing insulation. One of the most recognized frameworks in this domain is the Predicted Mean Vote (PMV) model, formulated by Fanger, which remains a cornerstone of thermal comfort assessment in engineering practice [1].

Despite its foundational importance, the PMV model and similar generalized approaches have faced increasing scrutiny. These models, while theoretically robust, are limited in their ability to adapt to individual variability and dynamic environmental contexts. For instance, they often fail to capture the subjective nuances of personal comfort, which can be influenced by cultural background, physiological states, and long-term exposure patterns. Empirical studies have shown substantial discrepancies between PMV predictions and actual occupant feedback in real-world settings [2], [3].

In response to these limitations, the field has seen a shift toward personalized thermal comfort modeling. These models aim to capture individual preferences and adapt over time through learning mechanisms that integrate occupant feedback and behavioral data. Personalized models are increasingly seen as critical components of smart building ecosystems, particularly in light of growing expectations for occupant-centric environmental control systems [4], [5].

This paradigm shift has been enabled by rapid technological advancements in the domains of the Internet of Things (IoT) and wearable sensing. IoT infrastructure in modern buildings allows for high-resolution, multi-modal data collection at scale. Environmental sensors embedded throughout indoor spaces continuously monitor parameters such as temperature, humidity, and CO_2 levels. Simultaneously, wearable devices enable noninvasive tracking of physiological signals like skin temperature, heart rate variability, and electrodermal activity—proxies that have demonstrated strong correlations with thermal sensation and comfort levels [6], [7].

These technologies support the development of data-driven comfort models that move beyond static rule-based systems toward dynamic, learning-based frameworks. As Boutahri and Tilioua highlight, integrating machine learning into thermal comfort prediction allows for systems that are both adaptive and energy-efficient, effectively balancing occupant well-being with sustainability goals [8]. Furthermore, the WELL Building Standard underscores the increasing institutional emphasis on health and well-being in architectural design, positioning thermal comfort as a core component of indoor environmental quality.

In summary, while classical models like PMV provide a valuable theoretical basis, the emergence of personalized, data-driven approaches—fueled by IoT and wearable technologies—represents a transformative development in smart environment research. The motivation for this thesis lies in exploring how these new data modalities and modeling techniques can be systematically leveraged to enhance the prediction and personalization of thermal comfort in real-world settings.

1.2 Problem Statement

Although substantial progress has been made in thermal comfort modeling, existing approaches face considerable challenges when applied to real-world environments. One of the most pressing limitations of current models—particularly those based on static physical principles like the Predicted Mean Vote (PMV) or the Adaptive Comfort Model—is their inability to accommodate individual variability and temporal dynamics in user preferences. These models often assume homogeneity across occupants and rely on fixed relationships between environmental inputs and comfort outcomes [1], [2].

However, thermal comfort is inherently personalized. Factors such as gender, age, activity level, clothing insulation, and even cultural background can significantly modulate comfort perceptions. For example, studies have documented consistent comfort preference differences between men and women, suggesting that generalized models systematically misrepresent subpopulations [9]. In addition to inter-individual variability, thermal preferences are also context-sensitive and time-dependent. Daily rhythms, seasonal changes, recent thermal exposure, and behavioral adaptations all influence thermal comfort in ways that static models cannot readily capture [10].

These challenges are magnified in smart building environments, where diverse occupants with varying schedules and preferences co-exist. Developing a single predictive model that generalizes well across users, locations, and time remains difficult, often resulting in suboptimal comfort control and user dissatisfaction. Moreover, user-specific models trained on one individual's data typically fail to transfer effectively to others or across different climate contexts. This hampers scalability and generalizability, particularly in systems where data collection is costly or privacy-sensitive [11], [12].

The need for flexible modeling approaches that can transfer knowledge between users and adapt to dynamic real-world conditions is thus evident. Addressing this problem requires methods that not only personalize thermal comfort prediction but also support cross-domain learning, enabling models to generalize across populations, environments, and temporal settings. This thesis tackles this gap by investigating the mechanisms and methodologies for transferable, user-centric thermal comfort modeling in smart environments.

1.3 Research Objectives

The central aim of this research is to advance personalized thermal comfort modeling by exploring methods for user preference transfer and model adaptability in smart environments. This involves integrating techniques from machine learning, domain adaptation, and spatiotemporal modeling to enhance comfort prediction and generalization capabilities.

The specific research objectives are as follows:

- Explore transferability of user preferences in real-world settings. Investigate the feasibility and limitations of transferring thermal comfort models across individuals, buildings, and climates, with the goal of minimizing the need for user-specific data collection.
- Develop and evaluate model adaptation techniques.

Apply and assess domain adaptation strategies—specifically *feature alignment* using Maximum Mean Discrepancy (MMD) and *instance reweighting* based on sample importance—to enable models to adjust to target domains with different data distributions.

- Implement spatiotemporal modeling with CNN-LSTM architectures. Design and analyze hybrid models that leverage Convolutional Neural Networks (CNNs) for local spatial feature extraction and Long Short-Term Memory (LSTM) networks for capturing temporal dynamics in user-environment interactions.
- Evaluate model performance across metrics and contexts. Conduct experiments to assess the effectiveness of proposed methods using realworld datasets, considering classification accuracy, adaptability, and cross-user generalization performance.

Collectively, these objectives support the development of more intelligent, user-aware, and transferable thermal comfort systems, suitable for deployment in sensor-rich smart environments.

1.4 Scope and Limitations

This thesis is centered on advancing data-driven thermal comfort modeling in smart environments by integrating user-centered machine learning techniques and domain adaptation frameworks. The scope of the work is both technically and conceptually delimited to ensure depth of investigation within a manageable experimental framework.

Scope of the Study

• Focus on multimodal data integration.

The primary emphasis is placed on the fusion of environmental and physiological data streams for personalized comfort modeling. Input sources include temperature, humidity, CO_2 levels, air velocity, and physiological indicators such as skin temperature and heart rate variability, when available. These variables are modeled as multivariate time series to capture dynamic interactions between the user and their ambient environment.

• Model-centric exploration of personalization.

The research focuses on algorithmic strategies for adapting thermal comfort models to individual users. Particular attention is given to methods that enable transfer of learned user preferences across contexts—geographical, architectural, or demographic—using techniques such as Maximum Mean Discrepancy (MMD) and instance reweighting.

• Use of pre-collected data.

The study leverages existing datasets collected from smart building environments and wearable sensors. Chief among these is the ASHRAE Global Thermal Comfort Database II [3], which provides a rich set of occupant-reported comfort responses alongside corresponding environmental measurements. All experiments and model evaluations are conducted offline, based solely on this pre-existing data without the use of synthetic simulations or emulated interventions.

Limitations of the Study

• No hardware development.

This thesis does not engage in the design or fabrication of sensing or actuation hardware. It assumes the availability of an IoT-enabled infrastructure that provides periodic sensor readings and optionally wearable data from users. Integration protocols, data transmission pipelines, and device calibration fall outside the research scope.

• Absence of live deployment trials.

Real-world system integration is discussed at the conceptual and architectural level. However, due to resource and time constraints, the study does not implement or validate models in live, large-scale deployments. Findings are instead based on controlled offline evaluation metrics and simulation-based inference.

• Limited generalization guarantees.

While the thesis investigates cross-domain transfer, it does so within the constraints of the datasets used. Generalizability to highly divergent environments (e.g., industrial settings, rural dwellings, extreme climates) cannot be definitively concluded from the results presented.

• No formal comfort validation experiments.

The assessment of thermal comfort relies on collected labels (e.g., Thermal Sensation Votes) and proxy features. The thesis does not include psychophysical experiments or longitudinal field studies to validate perceived comfort outcomes under proposed models.

By delineating the boundaries of the research, this section clarifies the methodological focus on personalized thermal comfort modeling from sensor data and ensures that conclusions are interpreted within a realistic and academically sound framework.

1.5 Thesis Structure

This thesis is organized into six main chapters, each progressively building upon the core research goals and empirical investigations.

Chapter 2 – Literature Review

This chapter presents a comprehensive review of the existing literature in smart environments and personalized thermal comfort modeling. It synthesizes findings from recent advancements in environmental sensing, adaptive control, and user-centric HVAC systems. The chapter serves to contextualize the research within ongoing trends and identifies critical knowledge gaps that this thesis seeks to address.

Chapter 3 – Theoretical Background

This chapter introduces the foundational theories of thermal comfort, including the widely adopted Predicted Mean Vote (PMV) and Predicted Percentage of Dissatisfied (PPD) models, the adaptive comfort model, and psychological approaches. It then transitions to core machine learning paradigms relevant to this study: supervised and unsupervised learning, reinforcement learning, and transfer learning. The chapter concludes with a detailed taxonomy of feature types and a discussion on time resolution, user-specific data modeling, and temporal embeddings.

Chapter 4 – Experiments and Results

This chapter outlines the experimental framework, beginning with an overview of the ASHRAE Global Thermal Comfort Database II. It then details exploratory data analysis, baseline model performance, and in-depth evaluations of machine learning algorithms for

comfort prediction, including logistic regression, random forests, multilayer perceptrons, and gradient boosting. A key focus is placed on user-level variability, feature relevance, and the performance implications of per-user adaptation. The chapter concludes with a dedicated section on transfer learning, exploring climate-adaptive model generalization across domains.

Chapter 5 – Discussion

This chapter reflects on the empirical findings, examining their implications for the personalization and deployment of thermal comfort models. It discusses model limitations, the impact of spatiotemporal dynamics, and the broader feasibility of integrating adaptive systems into intelligent building management. Special attention is given to the role of user preferences in driving model efficacy and energy-aware environmental adaptation.

Chapter 6 – Future Work

The final chapter outlines prospective research directions. It proposes the use of richer data modalities including physiological and contextual signals, advanced domain adaptation techniques, user clustering methods, lightweight models for edge deployment, hybrid physical-statistical models, active learning for reducing user feedback burden, and strategies for end-to-end system integration within IoT-enhanced smart environments.

Together, these chapters provide a structured and in-depth investigation into the challenges and opportunities of user-adaptive comfort modeling in real-world smart building contexts.

Chapter 2

Literature Review

2.1 Foundations of Thermal Comfort Modeling

2.1.1 Physical Models: PMV and PPD

The Predicted Mean Vote (PMV) and Predicted Percentage of Dissatisfied (PPD) models, developed by Fanger in 1970, are seminal in the field of thermal comfort assessment [1]. The PMV model predicts the mean thermal sensation vote of a large group of people on a seven-point scale ranging from cold (-3) to hot (+3), based on the heat balance of the human body. The PPD model estimates the percentage of occupants likely to feel thermally dissatisfied in a given environment.

These models consider six primary factors: air temperature, mean radiant temperature, relative humidity, air velocity, metabolic rate, and clothing insulation. The PMV model is mathematically expressed as:

$$PMV = (0.303e^{-0.036M} + 0.028) \cdot [(M - W) - E_d - E_r - C - R]$$
(2.1)

where M is the metabolic rate, W is the external work, E_d is the heat loss through diffusion, E_r is the evaporative heat loss, C is the convective heat loss, and R is the radiative heat loss.

Despite their widespread adoption, these models have limitations. They were developed under steady-state conditions and may not accurately predict comfort in dynamic or naturally ventilated environments. Moreover, they assume uniformity among occupants, neglecting individual differences in thermal perception [13].

2.1.2 Adaptive Comfort Models

Recognizing the limitations of static models like PMV/PPD, researchers developed adaptive comfort models that account for occupants' behavioral, physiological, and psychological adaptations to their thermal environment. The adaptive model posits that acceptable indoor temperatures are influenced by outdoor climatic conditions and occupants' expectations [2].

Field studies, such as the ASHRAE RP-884 project, collected data from various buildings worldwide to support the adaptive model. The findings indicated that occupants in naturally ventilated buildings accept a wider range of temperatures compared to those in mechanically conditioned spaces. This adaptability is attributed to increased personal control over the environment and acclimatization to local climates [14].

The adaptive model has been incorporated into standards like ASHRAE 55 and EN 15251, providing guidelines for acceptable indoor temperatures based on outdoor conditions. However, its applicability is primarily limited to naturally ventilated buildings where occupants can exercise control over their environment.

2.1.3 Psychological and Behavioral Extensions

Beyond physical and adaptive models, psychological and behavioral factors significantly influence thermal comfort perceptions. Individual differences in thermal sensitivity, expectations, and control over the environment can lead to varying comfort levels among occupants [15].

Studies have shown that factors such as personal control, thermal history, and cultural expectations affect thermal comfort. For instance, occupants with greater control over their environment, such as operable windows or personal fans, report higher satisfaction levels. Moreover, thermal preferences can be shaped by past experiences and cultural norms, influencing occupants' comfort expectations [16].

Incorporating these psychological and behavioral aspects into thermal comfort models can enhance their predictive accuracy and relevance. Emerging research explores integrating physiological signals, such as heart rate variability, to develop personalized comfort models that account for individual differences [7].

2.2 Personalization and User-Centric Modeling

2.2.1 Personalized Comfort Models (PCMs)

Personalized Comfort Models (PCMs) represent a paradigm shift in thermal comfort research, moving beyond generalized models to account for individual differences in thermal perception. These models leverage user-specific features such as demographics, physiological signals, and behavioral patterns to predict thermal comfort more accurately.

Lee and Chun (2021) developed a thermal comfort prediction model utilizing physiological signals obtained from wearable devices. Their study demonstrated that variables like skin temperature and electrodermal activity significantly correlate with thermal states, achieving an 80% prediction accuracy using only physiological data [6]. Similarly, Jayathissa et al. (2020) introduced a methodology for collecting intensive longitudinal subjective feedback using smartwatches, enabling the development of individualized comfort models based on real-time data [17].

The motivation for personalization in energy-aware smart systems stems from the need to enhance occupant comfort while optimizing energy consumption. Traditional HVAC systems often fail to accommodate individual preferences, leading to energy inefficiencies and occupant discomfort. PCMs address this by enabling adaptive control strategies that respond to individual thermal needs, thereby improving both comfort and energy efficiency.

2.2.2 Challenges in Generalization

While PCMs offer significant advantages, they also present challenges, particularly concerning generalization across diverse populations. Thermal preferences vary widely among individuals due to factors such as age, gender, metabolic rate, and acclimatization. This variability complicates the development of models that can generalize effectively across different user groups.

Studies have highlighted the limitations of one-size-fits-all models. For instance, research indicates that seniors often prefer warmer temperatures compared to younger adults, reflecting age-related differences in thermal sensitivity [18]. Additionally, cultural and climatic backgrounds influence thermal comfort perceptions, further complicating model generalization.

To address these challenges, researchers have explored cohort-based modeling approaches. Quintana et al. (2022) proposed Cohort Comfort Models that group occupants based on similarity in thermal preferences, enabling more accurate predictions with less individual data [19]. This approach balances the need for personalization with the practical constraints of data collection.

2.2.3 Longitudinal Studies and Feedback Systems

Longitudinal studies play a crucial role in understanding and modeling thermal comfort over time. By collecting data across extended periods, researchers can capture temporal variations in thermal preferences and adapt models accordingly.

Tekler et al. (2023) introduced an active learning framework for personalized thermal comfort modeling, which iteratively updates the model based on occupant feedback. This approach reduces the data collection burden while maintaining high prediction accuracy, facilitating practical implementation in real-world settings [20].

Similarly, Gnecco et al. (2023) conducted long-term thermal comfort monitoring using wearable sensing techniques. Their study found significant correlations between environmental metrics and subjective perceptions, underscoring the value of continuous monitoring for adaptive comfort modeling [21].

These feedback systems enable dynamic adjustment of environmental controls, aligning with occupants' evolving comfort needs and contributing to energy-efficient building operations.

2.3 Machine Learning Approaches in Comfort Modeling

2.3.1 Supervised Learning Techniques

Supervised learning techniques have been extensively employed to predict thermal comfort levels by mapping input features to comfort indices. Common algorithms include logistic regression, decision trees, artificial neural networks (ANNs), and ensemble models such as Random Forest (RF) and Extreme Gradient Boosting (XGBoost).

Boutahri and Tilioua (2024) developed a predictive model utilizing Support Vector Machine (SVM), ANN, RF, and XGBoost to forecast the Predicted Mean Vote (PMV) index. Their study demonstrated that RF and XGBoost achieved superior performance, with accuracies of 96.7% and 96.4% respectively, highlighting the efficacy of ensemble methods in thermal comfort prediction [8].

Zhao et al. (2025) applied machine learning algorithms to predict thermal comfort levels in various building types and climates. Their findings emphasized the importance of incorporating macro-contextual variables, such as climate class and ventilation strategy, to enhance model accuracy [22].

2.3.2 Unsupervised Learning and Clustering

Unsupervised learning methods, including Principal Component Analysis (PCA), k-means clustering, and Density-Based Spatial Clustering of Applications with Noise (DBSCAN), have been utilized to uncover latent structures in thermal comfort data. These techniques facilitate feature reduction and user segmentation, enabling the identification of distinct comfort preference groups.

For instance, clustering algorithms have been employed to group occupants based on similar thermal preferences, aiding in the development of cohort-based comfort models. This approach allows for more tailored HVAC control strategies without the need for extensive individual data collection [19].

PCA has been instrumental in reducing the dimensionality of complex datasets, retaining the most significant features influencing thermal comfort. This simplification enhances the efficiency and interpretability of subsequent predictive models.

2.3.3 Temporal Models and Deep Architectures

Temporal models, particularly those based on deep learning architectures, have shown promise in capturing the dynamic nature of thermal comfort perceptions. Convolutional Neural Networks (CNNs) and Long Short-Term Memory (LSTM) networks are adept at modeling temporal dependencies in sequential data.

Somu et al. (2021) introduced a hybrid model combining CNN and LSTM architectures to predict thermal comfort by analyzing spatiotemporal relationships in physiological and environmental data. Their approach demonstrated improved accuracy in capturing individual thermal sensations over time [23].

2.3.4 Reinforcement and Active Learning

Reinforcement Learning (RL) and Active Learning (AL) methodologies have been applied to develop adaptive thermal comfort models that learn optimal control strategies through interaction with the environment and occupants.

Gao et al. (2019) proposed a Deep Reinforcement Learning framework for HVAC control, formulating the problem as a cost-minimization task balancing energy consumption and occupant comfort. Their model demonstrated the ability to adapt to varying conditions and preferences, leading to enhanced energy efficiency and comfort levels [24].

Tekler et al. (2023) introduced an Active Learning approach to personalize thermal comfort models, reducing the data collection burden by selectively querying the most informative data points. This strategy achieved high prediction accuracy with significantly less labeled data, facilitating scalable deployment in real-world settings [20].

2.4 Transfer Learning and Domain Adaptation

2.4.1 Domain Transfer Challenges

Transfer learning and domain adaptation have emerged as pivotal techniques in thermal comfort modeling, particularly when addressing the variability across different environments and populations. However, several challenges impede the seamless transfer of models across domains.

Yang et al. (2025) highlighted the limited generalizability of thermal comfort models when applied to diverse climatic conditions and building types. Their study demonstrated that models trained in one domain often underperform when deployed in another, primarily due to discrepancies in environmental variables and occupant behaviors [12].

Similarly, Gao et al. (2021) investigated the application of transfer learning for thermal comfort prediction across multiple cities. Their findings underscored the necessity of domain adaptation techniques to mitigate the performance degradation caused by domain shifts [11].

2.4.2 Adaptation Techniques

To address the challenges of domain shifts, various adaptation techniques have been proposed. Instance reweighting adjusts the importance of source domain samples to better align with the target domain distribution. This method has been effective in enhancing model performance in new environments [25].

Adversarial training introduces a domain discriminator to encourage the model to learn domain-invariant features. Ganin et al. (2015) proposed the Domain-Adversarial Neural Network (DANN), which has been adapted for thermal comfort prediction to improve cross-domain generalization [26].

Maximum Mean Discrepancy (MMD)-based alignment measures the distance between source and target domain distributions in a reproducing kernel Hilbert space. By minimizing this distance, models can achieve better alignment and improved performance in the target domain [27].

2.4.3 Hybrid Models

Hybrid models that combine physics-based and data-driven approaches offer a promising avenue for thermal comfort prediction. Zhou et al. (2021) developed a hybrid model integrating the Predicted Mean Vote (PMV) index with machine learning techniques to enhance prediction accuracy. Their approach leverages the strengths of both methodologies, resulting in improved performance across various conditions [28]. Such hybrid models not only improve predictive capabilities but also provide interpretability, which is crucial for practical applications in building management systems.

2.5 Integration into Smart Systems

2.5.1 IoT and Wearables

The integration of Internet of Things (IoT) devices and wearable technologies has significantly advanced the continuous monitoring of environmental and physiological parameters relevant to thermal comfort. Wearable sensors capable of measuring skin temperature, heart rate, and heart rate variability (HRV) provide real-time data that can be utilized to develop personalized thermal comfort models [29].

The WELL Building Standard emphasizes the importance of monitoring thermal comfort parameters, advocating for the use of sensors to continuously assess environmental conditions such as temperature, humidity, and CO_2 levels. This continuous monitoring enables building managers and occupants to make informed decisions to maintain optimal thermal comfort [30].

Liu et al. (2019) demonstrated the efficacy of using wearable sensors to develop personal thermal comfort models. Their study involved collecting physiological data from participants over extended periods, resulting in models that achieved a median prediction accuracy of 78% [29].

2.5.2 Edge Deployment and Privacy

Deploying thermal comfort models on edge devices presents opportunities for real-time data processing and enhanced privacy. Edge computing minimizes latency and reduces the reliance on cloud-based services, thereby mitigating potential privacy concerns associated with data transmission.

Model compression techniques such as pruning and quantization are essential for facilitating the deployment of complex models on resource-constrained edge devices. Liang et al. (2021) provided a comprehensive survey on these techniques, highlighting their effectiveness in reducing model size and computational requirements without significantly compromising accuracy [31].

Recent advancements have demonstrated the feasibility of deploying compressed models on edge devices with high accuracy and low inference times. For instance, a study achieved a 92.5% accuracy with an inference time of 20 ms by employing structured pruning and dynamic quantization techniques [32].

2.5.3 Interface Design and User Interaction

The design of human-computer interfaces (HCI) plays a crucial role in the effectiveness of personalized environmental control systems. User-friendly interfaces facilitate occupant engagement, allowing individuals to provide feedback and adjust settings to their comfort preferences.

Zhu et al. (2021) explored adaptive HCI within the context of Industry 5.0, emphasizing the need for interfaces that can adapt to user behaviors and preferences. Their work underscores the importance of designing interfaces that are intuitive and responsive to enhance user satisfaction and system efficiency [33].

Effective HCI design in thermal comfort systems should prioritize usability, accessibility, and responsiveness. Incorporating features such as real-time feedback, customizable settings, and clear visualizations can empower users to actively participate in managing their thermal environment.
Chapter 3

Theoretical Background

3.1 Foundations of Thermal Comfort Modeling

3.1.1 Predicted Mean Vote (PMV) and Predicted Percentage of Dissatisfied (PPD)

The Predicted Mean Vote (PMV) model, introduced by Fanger [1], is designed to predict the average thermal sensation of a group of people under uniform environmental conditions. It is based on the ASHRAE 7-point thermal sensation scale:

-3 Cold

-2 Cool

-1 Slightly cool

0 Neutral

+1 Slightly warm

+2 Warm

+3 Hot

The PMV index is calculated using the equation:

$$PMV = \left[0.303 \, e^{-0.036M} + 0.028\right] \cdot L$$

where:

- M: metabolic rate $[W/m^2]$
- L: thermal load the net heat that the human body must lose to maintain thermal equilibrium

The thermal load L is defined by:

$$L = M - W - E_{\rm res} - C_{\rm res} - R - C - E_d - E_s$$

This expression captures the balance between internal heat production and the mechanisms of heat dissipation:

- M: energy generated by metabolic activity
- W: mechanical work done by the body (typically negligible indoors)
- $E_{\rm res}$: latent heat loss through evaporation in respiration
- $C_{\rm res}$: sensible heat loss from exhaled air
- R: radiative heat loss from the skin to cooler surrounding surfaces
- C: convective heat loss due to air movement
- E_d : insensible perspiration (diffusion through the skin)
- E_s : sensible perspiration (active sweating)

Essentially, L reflects the excess or deficit of heat in the body. A positive L indicates that the body is too warm (requiring cooling), while a negative L indicates a heat deficit (requiring warming). The PMV then interprets this thermal imbalance as a sensation vote on the ASHRAE scale.

The associated dissatisfaction metric is the Predicted Percentage of Dissatisfied (PPD):

$$PPD = 100 - 95 e^{-0.03353 \cdot PMV^4 - 0.2179 \cdot PMV^2}$$

This equation empirically estimates the percentage of occupants likely to feel thermally uncomfortable under given conditions. Even at optimal comfort (PMV = 0), about 5% of individuals may remain dissatisfied due to inter-individual variability [5].

3.1.2 Adaptive Comfort Model

The adaptive comfort model, developed by de Dear and Brager [2], considers behavioral and psychological adaptations to outdoor climate. It defines a temperature range in naturally ventilated buildings that shifts according to the outdoor thermal history:

$$T_{op}^{neutral} = 0.31 \cdot T_{rm} + 17.8$$

with the running mean temperature T_{rm} defined as:

$$T_{rm} = (1 - \alpha) \sum_{i=1}^{n} \alpha^{i-1} T_{od,i}$$

where $T_{od,i}$ is the mean outdoor temperature *i* days ago, and $\alpha \approx 0.8$. This model allows for broader comfort thresholds, especially where occupants can adjust clothing, open windows, or modify their environment.

3.1.3 Data-Driven Comfort Models

In contrast to physics-based models, data-driven approaches leverage machine learning to estimate comfort from observational data:

$$y = f(\mathbf{x}) + \epsilon$$

Here:

- **x**: input feature vector (e.g., air temperature, humidity, clothing, time)
- y: observed comfort feedback (e.g., ASHRAE scale score)
- ϵ : residual error due to unmeasured or subjective factors

Model training involves minimizing a loss function. In regression settings, the *Mean* Squared Error (MSE) is common:

$$\mathcal{L}_{\text{MSE}} = \frac{1}{n} \sum_{i=1}^{n} (\hat{y}_i - y_i)^2$$

where:

- \hat{y}_i : model prediction for sample *i*
- y_i : ground truth label
- *n*: number of training examples

For classification models (e.g., predicting discrete ASHRAE classes), the output is a probability distribution over comfort classes:

$$\hat{\mathbf{y}}_i = [\hat{y}_{i,1}, \dots, \hat{y}_{i,C}], \quad \sum_{c=1}^C \hat{y}_{i,c} = 1$$

The Cross-Entropy (CE) loss function measures the divergence between predicted and true class distributions:

$$\mathcal{L}_{\text{CE}} = -\frac{1}{n} \sum_{i=1}^{n} \sum_{c=1}^{C} y_{i,c} \log(\hat{y}_{i,c})$$

where:

- $y_{i,c}$: one-hot encoded true label
- $\hat{y}_{i,c}$: predicted probability for class c
- C: number of classes

These data-driven models enable dynamic personalization and are particularly useful in real-time applications such as personal environmental control systems [8], [34].

3.1.4 Cognitive and Psychological Models of Comfort

Comfort perception is also influenced by psychological factors. A heuristic formulation is:

$$T_{\text{neutral}} = T_{\text{phys}} + \phi(\bar{T}_{\text{past}}, \text{control}, \text{mood})$$

where:

- $T_{\rm phys}$: neutral temperature from physiological model
- \bar{T}_{past} : average of previously experienced temperatures
- control: perceived ability to adjust the environment
- mood: transient emotional state

Research shows that the perception of control—such as access to windows or thermostats—can improve thermal satisfaction even when actual conditions remain unchanged [5], [9]. Such models support human-in-the-loop strategies and adaptive interface design in smart environments.

3.2 Machine Learning Fundamentals

3.2.1 Introduction to Machine Learning

Machine Learning (ML) is a branch of artificial intelligence (AI) that focuses on the development of algorithms and statistical models enabling computers to perform specific tasks without explicit instructions. Unlike traditional programming, where rules are hard-coded, ML systems learn patterns from data, allowing them to make decisions or predictions based on new inputs. This section provides an overview of the foundational concepts, types, training processes, applications, and challenges associated with machine learning.

3.2.2 Principles of Machine Learning

At its core, machine learning involves the following key principles:

- **Data Representation**: Data is structured into a format suitable for analysis, often as feature vectors where each feature represents a measurable attribute of the phenomenon under study.
- **Modeling**: Algorithms are employed to create models that capture patterns and relationships within the data, facilitating predictions or decisions.
- **Training**: Models are trained using datasets, adjusting internal parameters to minimize errors between predicted and actual outcomes.

- **Evaluation**: Trained models are assessed using separate validation datasets to ensure they generalize well to unseen data. Metrics such as accuracy, precision, recall, and F1-score are commonly used.
- **Prediction**: Once validated, models can make predictions or decisions on new data, applying learned patterns to real-world scenarios.

3.2.3 Types of Machine Learning

Machine learning encompasses several paradigms, each suited to different types of problems:

Supervised Learning

In supervised learning, models are trained on labeled datasets, where each input is paired with a known output. The objective is to learn a mapping from inputs to outputs, enabling the prediction of outcomes for new, unseen data.

Common Algorithms: Linear Regression, Logistic Regression, Support Vector Machines (SVM), Decision Trees, Random Forests, Neural Networks.

Applications: Image classification, spam detection, medical diagnosis, speech recognition.

Unsupervised Learning

Unsupervised learning deals with unlabeled data, aiming to uncover hidden structures or patterns without predefined outputs.

Common Algorithms: K-Means Clustering, Hierarchical Clustering, Principal Component Analysis (PCA), t-Distributed Stochastic Neighbor Embedding (t-SNE).

Applications: Customer segmentation, anomaly detection, market basket analysis, gene expression analysis.

Reinforcement Learning

Reinforcement learning involves training agents to make a sequence of decisions by rewarding or penalizing actions, with the goal of maximizing cumulative rewards over time. **Common Algorithms**: Q-Learning, Deep Q-Networks (DQN), Policy Gradient Methods.

Applications: Robotics, game playing (e.g., AlphaGo), autonomous driving, resource management.

3.2.4 Training Machine Learning Models

The process of training machine learning models involves several critical steps:

1. **Data Collection**: Gathering relevant and diverse data is fundamental, as the quality and quantity of data directly influence model performance.

- 2. **Data Preprocessing**: Raw data often contains noise or inconsistencies. Preprocessing includes cleaning data, handling missing values, and normalizing or standardizing features.
- 3. Feature Engineering: Selecting and transforming variables to enhance model learning capabilities. Effective feature engineering can significantly improve model accuracy.
- 4. **Model Selection**: Choosing appropriate algorithms based on the problem type and data characteristics. Different algorithms have varying strengths and are suited to specific tasks.
- 5. **Training**: Adjusting model parameters using training data to minimize prediction errors. Techniques like gradient descent are commonly used for optimization.
- 6. Validation: Evaluating model performance on a separate validation set to finetune hyperparameters and prevent overfitting.
- 7. **Testing**: Assessing the final model on an independent test set to estimate its performance on unseen data.

3.2.5 Applications of Machine Learning

Machine learning has a broad spectrum of applications across various industries:

- **Healthcare**: Disease prediction, medical imaging analysis, personalized treatment plans, drug discovery.
- **Finance**: Fraud detection, stock market prediction, credit scoring, algorithmic trading.
- **Retail**: Demand forecasting, inventory management, recommendation systems, customer segmentation.
- **Transportation**: Autonomous vehicles, traffic prediction, route optimization, ridesharing services.
- Natural Language Processing (NLP): Language translation, sentiment analysis, chatbots, speech recognition.
- **Computer Vision**: Object detection, facial recognition, image classification, video analysis.
- **Entertainment**: Content recommendation, personalized advertising, user behavior analysis.

3.2.6 Challenges and Future Directions

Despite significant advancements, machine learning faces several challenges:

- **Data Quality and Quantity**: Obtaining high-quality, labeled data is often resourceintensive, yet essential for effective model training.
- Interpretability: Complex models, especially deep learning networks, can act as "black boxes," making it difficult to understand decision-making processes.
- **Bias and Fairness**: Models may inherit biases present in training data, leading to unfair or discriminatory outcomes.
- **Scalability**: Training models on large datasets requires substantial computational resources, posing scalability issues.
- Security and Privacy: Ensuring data privacy and security is crucial, particularly in sensitive domains like healthcare and finance.

Looking ahead, advancements in explainable AI, federated learning, and quantum computing hold promise for addressing these challenges, paving the way for more transparent, efficient, and secure machine learning applications.

3.3 Machine Learning and Data-Driven Techniques

The advent of machine learning (ML) has revolutionized the modeling of thermal comfort by enabling data-driven approaches that adapt to individual preferences and environmental variations. This section delves into various ML techniques applicable to thermal comfort prediction and control.

3.3.1 Supervised Learning

Supervised learning involves training a model on a labeled dataset, where the input features $\mathbf{X} = \{x_1, x_2, \ldots, x_n\}$ are associated with known outputs $\mathbf{y} = \{y_1, y_2, \ldots, y_n\}$. The goal is to learn a mapping function $f : \mathbf{X} \to \mathbf{y}$ that can predict outputs for unseen inputs.

Logistic Regression

Logistic regression is a statistical model used for binary classification problems. It estimates the probability that a given input \mathbf{x} belongs to a particular class. The model is defined as:

$$P(y = 1 | \mathbf{x}) = \frac{1}{1 + e^{-(\mathbf{w}^{\top} \mathbf{x} + b)}}$$
(3.1)

where:

• **w** is the weight vector,

- *b* is the bias term,
- *e* is the base of the natural logarithm.

The model is trained by minimizing the binary cross-entropy loss:

$$\mathcal{L} = -\frac{1}{n} \sum_{i=1}^{n} [y_i \log(\hat{y}_i) + (1 - y_i) \log(1 - \hat{y}_i)]$$
(3.2)

where \hat{y}_i is the predicted probability for sample *i*.

Decision Trees

Decision trees are flowchart-like structures where internal nodes represent tests on features, branches represent outcomes of the tests, and leaf nodes represent class labels or regression values. The tree is built by recursively partitioning the data to minimize impurity measures such as Gini impurity or entropy.

Random Forests

Random forests are ensemble models that construct multiple decision trees during training and output the mode of the classes (classification) or mean prediction (regression) of the individual trees. The model reduces overfitting by averaging multiple deep decision trees trained on different parts of the same dataset.

Extreme Gradient Boosting (XGBoost)

XGBoost is an optimized distributed gradient boosting library designed to be highly efficient and flexible. It builds models in a stage-wise fashion and generalizes them by allowing optimization of an arbitrary differentiable loss function. The objective function is:

$$\mathcal{L}(\phi) = \sum_{i=1}^{n} l(y_i, \hat{y}_i) + \sum_{k=1}^{K} \Omega(f_k)$$
(3.3)

where:

- *l* is a differentiable convex loss function,
- $\Omega(f_k) = \gamma T + \frac{1}{2}\lambda ||w||^2$ is the regularization term,
- T is the number of leaves in the tree,
- w is the vector of scores on leaves.

Multilayer Perceptron (MLP)

An MLP is a class of feedforward artificial neural network that consists of at least three layers of nodes: an input layer, a hidden layer, and an output layer. Each node (except for the input nodes) is a neuron that uses a nonlinear activation function. The output of an MLP is given by:

$$\hat{y} = f(\mathbf{W}_2 \cdot g(\mathbf{W}_1 \cdot \mathbf{x} + \mathbf{b}_1) + \mathbf{b}_2) \tag{3.4}$$

where:

- $\mathbf{W}_1, \mathbf{W}_2$ are weight matrices,
- \mathbf{b}_1 , \mathbf{b}_2 are bias vectors,
- g is the activation function (e.g., ReLU, sigmoid),
- f is the output activation function (e.g., softmax for classification).

3.3.2 Unsupervised Learning and Dimensionality Reduction

Unsupervised learning deals with unlabeled data, aiming to uncover hidden structures or patterns.

Principal Component Analysis (PCA)

PCA is a dimensionality reduction technique that transforms the data into a new coordinate system such that the greatest variance lies on the first principal component, the second greatest variance on the second component, and so on. Mathematically, PCA solves the eigenvalue decomposition of the data covariance matrix.

K-Means Clustering

K-Means clustering partitions n observations into k clusters in which each observation belongs to the cluster with the nearest mean. The objective is to minimize the withincluster sum of squares (WCSS):

$$\arg\min_{S} \sum_{i=1}^{k} \sum_{\mathbf{x} \in S_i} \|\mathbf{x} - \mu_i\|^2$$
(3.5)

where:

- $S = \{S_1, S_2, \dots, S_k\}$ are the clusters,
- μ_i is the centroid of cluster S_i .

DBSCAN (Density-Based Spatial Clustering of Applications with Noise)

DBSCAN is a density-based clustering algorithm that groups together points that are closely packed together, marking as outliers points that lie alone in low-density regions. It requires two parameters: ε (maximum radius of the neighborhood) and *MinPts* (minimum number of points required to form a dense region).

3.3.3 Reinforcement Learning

Reinforcement learning (RL) is an area of machine learning concerned with how agents ought to take actions in an environment to maximize cumulative reward. An RL problem is typically modeled as a Markov Decision Process (MDP) defined by a tuple (S, A, P, R, γ) , where:

- S is the set of states,
- A is the set of actions,
- *P* is the state transition probability matrix,
- *R* is the reward function,
- γ is the discount factor.

The goal is to find a policy $\pi: S \to A$ that maximizes the expected return:

$$\mathbb{E}\left[\sum_{t=0}^{\infty} \gamma^t R(s_t, a_t)\right] \tag{3.6}$$

3.3.4 Transfer Learning

Transfer learning refers to the process of improving learning performance in a target domain by leveraging knowledge acquired from a related source domain. This approach is particularly relevant in thermal comfort modeling, where collecting sufficient labeled data in each new building, climate, or demographic group can be both expensive and timeconsuming. Transfer learning enables the reuse of models or representations developed in one setting (e.g., a specific climate zone or building type) for another, accelerating deployment and improving prediction accuracy in data-sparse contexts.

Motivations and Challenges

In the context of smart buildings and personalized comfort systems, one frequently encounters variability in environmental conditions, occupant preferences, sensor configurations, and HVAC control strategies across different deployment environments. These variations introduce *domain shift*, where the statistical properties of the input data differ between source and target environments. For example, a thermal comfort model trained on data from a temperate climate may perform poorly in a tropical setting due to differences in indoor-outdoor gradients, humidity profiles, and occupant clothing habits. Formally, we denote the source domain as:

 $\mathcal{D}_S = \{X_S, P_S(X)\}, \text{ and the target domain as: } \mathcal{D}_T = \{X_T, P_T(X)\}$

where X is the input feature space and P(X) the marginal distribution. The goal of transfer learning is to approximate the conditional distribution $P_T(Y|X)$ in the target domain using knowledge from $P_S(Y|X)$, despite the divergence in input distributions $P_S(X) \neq P_T(X)$.

Key challenges include:

- Overfitting to the source domain's specific environmental features.
- Misalignment of feature semantics across domains (e.g., sensor drift or scale variation).
- Scarcity of labeled data in the target domain for fine-tuning or validation.

Taxonomy of Transfer Learning Strategies

Transfer learning techniques can be categorized as:

- Inductive Transfer Learning: $P_S(X) \approx P_T(X)$, but $P_S(Y|X) \neq P_T(Y|X)$. Common in personalization.
- Transductive Transfer Learning: $P_S(X) \neq P_T(X)$, but $P_S(Y|X) = P_T(Y|X)$. Used across climates or buildings.
- Unsupervised Transfer Learning: Both source and target domains lack labels. Used for latent representation learning.

Pretraining and Fine-Tuning Architectures

A common approach involves:

- 1. **Pretraining** on source domain \mathcal{D}_S to learn parameters θ via supervised learning.
- 2. Fine-tuning on target domain \mathcal{D}_T using limited labels, either with all parameters or partial adaptation (e.g., freezing early layers).

In CNN-LSTM architectures for thermal comfort, convolutional layers often extract spatial features while LSTM layers capture temporal dynamics. Fine-tuning enables these models to adapt to new occupancy schedules, building layouts, or sensor arrays.

Advanced Domain Adaptation Strategies

Several strategies can reduce domain discrepancy:

Feature Alignment Feature alignment aims to minimize the discrepancy between the distributions of source and target domain features by projecting both into a common, high-dimensional space. This is commonly achieved using Maximum Mean Discrepancy (MMD), which quantifies the distance between the mean embeddings of the two domains. The MMD is defined as:

$$\mathrm{MMD}(\mathcal{D}_S, \mathcal{D}_T) = \left\| \frac{1}{n_S} \sum_{i=1}^{n_S} \phi(x_i^S) - \frac{1}{n_T} \sum_{j=1}^{n_T} \phi(x_j^T) \right\|_{\mathcal{H}}^2$$

where ϕ is a feature mapping function into a reproducing kernel Hilbert space (RKHS) \mathcal{H} . This function transforms raw input features (e.g., temperature and humidity) into a space where similarity can be measured via inner products. A key property of RKHS is that it allows comparisons using kernel functions such as the Radial Basis Function (RBF), polynomial, or linear kernels, which enable implicit computation of high-dimensional distances without explicitly transforming the inputs.

Minimizing the MMD during training aligns the source and target distributions in \mathcal{H} , making them statistically similar. This term is typically added to the standard prediction loss in the overall objective:

$$\mathcal{L}_{\text{total}} = \mathcal{L}_{\text{pred}} + \lambda \cdot \text{MMD}(\mathcal{D}_S, \mathcal{D}_T)$$

where λ balances prediction accuracy and domain alignment. By reducing MMD, the model learns features that are domain-invariant and thus generalize better across different environmental contexts.

Instance Reweighting Instance reweighting addresses the domain shift problem by assigning different importance weights to source domain samples based on their similarity to the target domain distribution. Instead of treating all source samples equally, this approach increases the influence of samples that are more representative of the target domain. The reweighted loss function is expressed as:

$$\mathcal{L}_{\text{adapt}} = \sum_{i=1}^{n_S} \beta_i \cdot \ell(f(x_i^S), y_i^S)$$

where β_i is the importance weight for the *i*-th source sample, and ℓ is a loss function such as cross-entropy or mean squared error.

The weights β_i are typically estimated as the ratio of the target to source probability densities at each point:

$$\beta_i = \frac{P_T(x_i^S)}{P_S(x_i^S)}$$

This formulation prioritizes source instances that lie in regions of the feature space where the target distribution has high density. In practice, density ratios can be estimated using kernel density estimation or probabilistic classifiers trained to distinguish between source and target instances.

The reweighted loss \mathcal{L}_{adapt} is then used either as the main training loss or combined with other objectives (e.g., regularization or domain discrepancy terms) to improve generalization across domains. This method helps mitigate overfitting to irrelevant source patterns and enhances the model's transferability.

CNN-LSTM Architectures for Spatiotemporal Adaptation

The combination of Convolutional Neural Networks (CNNs) and Long Short-Term Memory (LSTM) networks forms a powerful hybrid architecture well-suited for thermal comfort prediction using sensor-rich, time-dependent data. CNNs excel at extracting spatial and local patterns from multivariate input features such as temperature, humidity, and CO_2 concentration. LSTMs are designed to capture temporal dependencies, enabling the model to learn from historical sequences of environmental and physiological states.

1D Convolutional Networks In the context of thermal comfort, the spatial structure typically refers to the feature channels collected at each time step, such as readings from environmental sensors or wearable devices. Therefore, 1D convolutions are applied along the temporal axis of these feature vectors:

$$\mathbf{z}_t = \text{CNN}(\mathbf{x}_{t-k:t})$$

Here, $\mathbf{x}_{t-k:t}$ represents a sliding window of feature vectors over k past time steps. Each convolutional filter learns local patterns such as increasing humidity, sudden temperature drops, or changes in heart rate, which may signal discomfort or behavioral adaptation.

Convolution Parameters The key hyperparameters in CNNs include:

- Kernel size: Determines the width of the convolutional window. A kernel size of 3 or 5 allows capturing local temporal transitions in the features.
- Stride: Controls the step size for moving the kernel along the input. A stride of 1 maintains resolution, while a larger stride reduces sequence length and computation.
- **Padding:** Zero-padding preserves the input length, enabling boundary information retention.
- Activation functions: Non-linear functions such as $\operatorname{ReLU}(\operatorname{ReLU}(x) = \max(0, x))$ are applied after convolution to introduce non-linearity, allowing the network to learn complex feature interactions.
- **Pooling layers:** Max-pooling or average-pooling can follow convolutions to reduce dimensionality and capture dominant features. Pooling helps with robustness to small variations and accelerates training.

Long Short-Term Memory Networks After spatial feature extraction, the resulting sequence $\mathbf{z}_{1:T}$ is fed into an LSTM layer, which models temporal dependencies. LSTM units are designed to mitigate the vanishing gradient problem, maintaining memory over long sequences via gated mechanisms:

$\mathbf{f}_t = \sigma(\mathbf{W}_f \cdot [\mathbf{h}_{t-1}, \mathbf{z}_t] + \mathbf{b}_f)$	(forget gate)
$\mathbf{i}_t = \sigma(\mathbf{W}_i \cdot [\mathbf{h}_{t-1}, \mathbf{z}_t] + \mathbf{b}_i)$	(input gate)
$\tilde{\mathbf{c}}_t = anh(\mathbf{W}_c \cdot [\mathbf{h}_{t-1}, \mathbf{z}_t] + \mathbf{b}_c)$	(candidate state)
$\mathbf{c}_t = \mathbf{f}_t \odot \mathbf{c}_{t-1} + \mathbf{i}_t \odot \widetilde{\mathbf{c}}_t$	(cell state)
$\mathbf{o}_t = \sigma(\mathbf{W}_o \cdot [\mathbf{h}_{t-1}, \mathbf{z}_t] + \mathbf{b}_o)$	(output gate)
$\mathbf{h}_t = \mathbf{o}_t \odot \tanh(\mathbf{c}_t)$	(hidden state output)

Each gate uses the sigmoid activation σ , controlling how information flows through time. The memory cell \mathbf{c}_t preserves long-term trends, while \mathbf{h}_t serves as the dynamic representation used for prediction.

Output Layers and Prediction The final hidden state \mathbf{h}_T can be passed through a stack of fully connected (dense) layers, possibly with dropout for regularization. These layers perform non-linear transformations to refine the representation for classification (e.g., comfort vs discomfort) or regression (e.g., predicted thermal sensation vote, TSV).

Architectural Extensions Several enhancements can be integrated into the base CNN-LSTM architecture to increase expressiveness, interpretability, and performance:

- Bidirectional LSTMs (BiLSTM): Standard LSTMs process input sequences from past to present. BiLSTMs augment this by also processing the sequence in reverse. This is particularly useful when the prediction target (e.g., retrospective comfort feedback) depends on both past and future context. Formally, the hidden state becomes a concatenation $\mathbf{h}_t = [\overrightarrow{\mathbf{h}_t}; \overleftarrow{\mathbf{h}_t}]$, capturing richer temporal information.
- Attention Mechanisms: Attention modules assign varying importance to different time steps. This is achieved by computing alignment scores and forming a context vector as a weighted sum of LSTM outputs:

$$\mathbf{c}_T = \sum_{t=1}^T \alpha_t \mathbf{h}_t$$
 where $\alpha_t = \frac{\exp(e_t)}{\sum_{k=1}^T \exp(e_k)}$

where $e_t = \text{score}(\mathbf{h}_t)$ represents a learned relevance score. Attention enhances interpretability and allows the model to focus on key moments affecting comfort.

• **Residual Connections:** Deep CNN layers often suffer from vanishing gradients. Residual blocks, popularized by ResNet, address this by allowing the input to bypass convolutions via identity mappings:

$$\mathbf{z}_t^{\text{out}} = \mathbf{z}_t^{\text{in}} + \text{CNN}(\mathbf{z}_t^{\text{in}})$$

This stabilizes training and allows deeper architectures without degradation.

- Multi-Scale CNNs: Multiple convolution paths with different kernel sizes (e.g., 3, 5, 7) can be applied in parallel to extract features at various temporal scales. Outputs from these paths are concatenated before feeding into LSTMs, enabling the network to capture both fine-grained and broad environmental dynamics.
- **Temporal Convolutional Networks (TCNs):** As an alternative or complement to LSTMs, TCNs use dilated causal convolutions to model long-range dependencies while preserving sequence order. TCNs provide parallelism in training and often outperform RNNs on time series tasks.

These extensions enhance the model's ability to generalize across diverse comfort contexts and improve learning efficiency, especially when applied to real-time, multi-sensor environments.

This hybrid CNN-LSTM model, particularly with the described enhancements, offers a robust framework for capturing both the short-term variability and long-term patterns inherent in personalized comfort prediction from rich sensor data.

Applications and Evaluation

Transfer learning has been applied to:

- Personalize comfort models across seasons and buildings [5]
- Transfer between geographical regions using ASHRAE databases [3]
- Enable real-time learning in wearable-driven systems [35]

Evaluation metrics include:

- Root Mean Squared Error (RMSE) or Mean Absolute Error (MAE) on target data.
- Classification Accuracy or F1 Score for ASHRAE scale prediction.
- Transfer Gain: $\Delta RMSE = RMSE_{baseline} RMSE_{transfer}$
- Domain Classification Accuracy (for adversarial training)

3.4 Feature Representation and Engineering

Effective feature representation is foundational to thermal comfort modeling in intelligent environments. The heterogeneity of data from users, environmental sensors, and interaction logs necessitates a structured approach to capture both immediate conditions and longer-term behavioral patterns. This section organizes features into four conceptual categories: user-centric data, environmental features, contextual/interaction variables, and temporal resolution considerations. We also highlight methods for encoding historical user experience through embeddings and memory structures.

3.4.1 User-Centric Data

Personal characteristics heavily influence thermal perception and behavioral adaptation. Incorporating user-specific features enables personalized modeling and enhances generalizability across heterogeneous populations.

- Static Demographics: Variables such as age, gender, and body mass index (BMI) serve as coarse but effective descriptors of thermoregulatory differences among individuals. These features are particularly useful when personalization must begin without historical data.
- **Physiological Proxies**: Metrics such as heart rate, skin temperature, and estimated metabolic rate—often derived from wearables—provide real-time signals of the user's thermal state. These data points support dynamic adaptation in systems designed for reactive control.
- Behavioral Feedback and Preference Histories: Explicit comfort feedback (e.g., thermal sensation votes), behavioral traces (e.g., thermostat interactions), and intervention patterns offer insight into subjective preference dynamics. Over time, such records help distinguish between transient discomfort and systemic misalignment.
- **Time-Series Embeddings**: To efficiently encode complex user histories, sequences of sensor readings and user feedback can be transformed into compact, fixed-size vector representations using models such as LSTMs, Transformers, or variational autoencoders. These embeddings preserve temporal dynamics while reducing dimensionality and can be used as inputs for downstream prediction tasks. They support both individual personalization and cross-user clustering [36].
- Historical Aggregates and Temporal Decay: When full time series are not available or computationally feasible, summary statistics such as exponentially weighted moving averages or recency-weighted feedback counts can approximate temporal patterns in user comfort.

3.4.2 Environmental Features

Environmental sensing provides the physical context in which thermal comfort is experienced. Rich, continuous environmental data allows for responsive modeling and finegrained control.

• Indoor Ambient Conditions: Temperature, relative humidity, air velocity, and CO₂ levels are standard inputs collected through IoT devices or building management systems (BMS). In more sophisticated settings, mean radiant temperature (MRT) and radiant asymmetry may also be captured to refine thermal load estimates.

- **Derived Metrics**: Features such as temperature change rate $(\Delta T/\Delta t)$, comfort envelope deviations, and rolling averages help identify transitional states, environmental stability, or anomalies. These derived features improve sensitivity to short-term discomfort triggers.
- Outdoor and Building Context: External weather data (e.g., solar radiation, wind speed) and envelope characteristics (e.g., window orientation, insulation quality) influence thermal behavior and energy dynamics. They are essential for adapting comfort predictions across varying spatial zones or building types.

3.4.3 Contextual and Interaction Features

Temporal and behavioral context shapes thermal experience beyond static sensor readings. These features capture how comfort is mediated by routine, location, and user interaction.

- **Temporal Encoding**: Time of day, day of week, and seasonal cycles influence thermal expectations and behavior. These can be encoded using cyclical transformations (e.g., $\sin(2\pi t/P), \cos(2\pi t/P)$) to preserve continuity and improve model interpretability.
- **Control and Interaction States**: User interventions such as HVAC adjustments, window operations, and fan usage reveal proactive comfort strategies. These are often encoded as binary or categorical features and provide a direct link between comfort perception and environmental response.
- Occupancy and Activity: Inferred from motion sensors, wearable devices, or scheduling systems, these features contextualize sensor data and help disambiguate causes of discomfort (e.g., heat stress vs. physical exertion).

3.4.4 Time Resolution and Historical Modeling

Temporal resolution significantly affects model accuracy, responsiveness, and deployment feasibility.

- High-Resolution Data (1–2 min): Suitable for real-time systems using wearable devices and responsive HVAC control. Offers high fidelity but can introduce noise and increase computational load.
- Medium Resolution (5–15 min): Common in BMS deployments, balancing detail and processing overhead. Captures most user-environment interactions effectively.
- Low Resolution (hourly or daily): Useful for long-term comfort profiling, design simulations, and energy planning. May smooth over important short-term variations.

• Historical Modeling Approaches:

- Time-Series Embeddings: Enable compression of long behavioral and environmental histories into predictive feature vectors [36].
- *Exponential Decay Functions*: Weight recent events more heavily, simulating the fading influence of older experiences.
- *Stateful Memory*: Use LSTM-style internal states or reinforcement learning agents to track and update user context dynamically across interaction episodes.

Chapter 4 Experiments and Results

4.1 Dataset Overview

This study utilizes the ASHRAE Global Thermal Comfort Database II, a comprehensive, real-world dataset comprising environmental, physiological, and subjective comfort data collected from thousands of building occupants across various climates and countries. It includes over 100,000 records spanning numerous buildings, climate zones, and subject profiles, making it a robust source for evaluating personalized thermal comfort models.

The experimental foundation of this thesis is the **ASHRAE II Thermal Comfort Dataset**, curated by the Center for the Built Environment (CBE) at the University of California, Berkeley. This dataset consolidates thermal comfort records from diverse field studies conducted in naturally ventilated, mixed-mode, and air-conditioned build-ings worldwide. Unlike controlled laboratory datasets, ASHRAE II captures occupants' responses in real-world settings, supporting more ecologically valid and generalizable thermal comfort models.

Each observation represents a snapshot in time, capturing a subject's thermal sensation and associated contextual factors. These are structured across four principal dimensions:

- Environmental Variables: Indoor measurements include air temperature, globe and radiant temperatures, relative humidity, air velocity (at multiple heights), and binary flags for control states (e.g., window open/closed, fan usage). Outdoor data—either directly measured or inferred from meteorological sources—include outdoor temperature, relative humidity, and a 7-day running mean temperature for adaptive modeling.
- **Personal Characteristics**: Demographic features include age, gender, height, and weight. Physiological parameters such as metabolic rate and clothing insulation are provided, with some variables derived from averaged activity levels over various time intervals (e.g., 10, 30, 60 minutes).
- Subjective Responses: Participants rate their thermal comfort using the 7-point ASHRAE thermal sensation scale (3 = cold to +3 = hot), as well as the acceptability and preference scales for both temperature and air movement.

• Spatiotemporal and Metadata: Each entry includes timestamp data, building and subject IDs, geographic location (country, city, latitude/longitude), and categorical climate classification (e.g., humid subtropical, tropical savanna). Metadata flags indicate the quality and source of measurements, helping to ensure traceability and filtering reliability during preprocessing.

To support supervised learning, the dataset includes unique identifiers for each building_id and subject_id, facilitating longitudinal and personalized modeling. The inclusion of PMV/PPD and SET calculations enables comparison against traditional thermal comfort standards (e.g., ISO 7730, ASHRAE 55).

Overall, the ASHRAE II dataset's structure and breadth make it an ideal candidate for personalized comfort modeling and transfer learning experiments in diverse environmental contexts.

4.2 Exploratory Data Analysis

This section provides an analytical overview of the thermal comfort dataset, highlighting the relationships between variables, their statistical behavior, and their relevance to the prediction task. Through visualizations, correlation studies, and dimensionality analysis, this exploratory phase informed the feature selection strategy, justified modeling choices, and revealed both the potential and limitations of the dataset.

4.2.1 Target Variable: Thermal Sensation

The target variable for all supervised learning experiments in this thesis is the thermal_sensation value, recorded as part of the ASHRAE 7-point thermal sensation scale. It is defined as follows:

- -3: Cold
- -2: Cool
- -1: Slightly Cool
- 0: Neutral
- +1: Slightly Warm
- +2: Warm
- +3: Hot

This variable is particularly advantageous for modeling purposes due to its availability across the majority of records and its granularity, which allows for nuanced classification beyond binary or ternary categorizations. However, the use of a 7-class formulation introduces greater complexity and ambiguity in prediction, especially considering the subjective nature of thermal comfort, which is influenced by latent psychological and physiological factors not fully captured in the dataset.

As such, while it provides a more descriptive framework for characterizing individual comfort, this target also inherently limits predictive performance, particularly for general-purpose models that are not personalized or context-aware.

4.2.2 Feature Correlation Analysis

A comprehensive correlation matrix was generated for all numerical features. Several intuitive and expected relationships were confirmed:

- Strong positive correlations between various indoor temperature measurements (e.g., air, globe, radiant, operative).
- Negative correlations between outdoor temperature and clothing insulation, as colder weather prompts heavier clothing.
- Seasonal interactions, such as increased fan usage in summer and greater clothing insulation in winter.

To assess the predictive value of individual features, correlations with the target variable thermal_sensation were computed. The following insights emerged:

- Air temperature showed the highest positive correlation, affirming its role as the most influential environmental variable.
- **Height** and **weight** exhibited negligible correlations, justifying their exclusion from the modeling pipeline.
- **Seasonal encodings** (e.g., summer, winter) revealed moderate correlations, supporting the inclusion of temporal context.
- A negative correlation with **clothing insulation** likely reflects adaptive behavior rather than causal influence.
- Features like window, fan, and door states showed modest yet notable correlations, hinting at potential for actionable interventions.



Figure 4.1: Correlation matrix of numerical features



Figure 4.2: Correlation of selected features with thermal sensation

4.2.3 Principal Component Analysis (PCA)

Principal Component Analysis was conducted on a subset of the most representative and readily collectible features, including both environmental and user-related variables. The selected features were:

- Air temperature
- Relative humidity
- Outdoor temperature
- Outdoor relative humidity
- Metabolic rate
- Clothing insulation
- Air velocity
- Gender (encoded as gender_male)
- Season (encoded as season_winter)

This selection reflects variables that are both frequently available in real deployments and theoretically linked to thermal comfort.

The cumulative explained variance plot showed no distinct elbow, implying that no small subset of components captures the majority of variance. This indicates the presence of nuanced, high-dimensional relationships justifying the use of full feature sets in the modeling process.



Figure 4.3: Cumulative variance explained by PCA components

4.2.4 User-Level Record Distribution

A key goal of this study is to explore the viability of personalized thermal comfort modeling. To assess this, we analyzed the distribution of the number of data records available per user. This metric provides insights into whether the dataset includes enough temporal data per individual to support per-user model training or transfer learning strategies.

Figure 4.4 illustrates the number of subjects who have at least a given number of data records. The results show that over 50 individuals have more than 110 recorded samples. This finding is significant, as datasets of this size per user enable the development of reliable personalized models and make the application of user-specific fine-tuning techniques—such as transfer learning—both feasible and meaningful.

This distribution also informs our model design decisions. Users with a sufficient history of measurements can support more data-intensive approaches, while those with fewer records may benefit from models that leverage population-level trends or shared representations. Understanding the variability in record availability is therefore essential for balancing generalization and personalization in thermal comfort prediction.



Figure 4.4: Number of users with at least N records

4.2.5 Per-User Correlation Distribution Analysis

To further understand individual variability in thermal comfort, we conducted a peruser correlation analysis between selected environmental and physiological features and the target variable, *thermal sensation*. For each user with sufficient data (typically ≥ 10 records), we computed the Pearson correlation between each selected feature and the user's reported thermal sensation. This produced a distribution of correlation coefficients for each feature across the user population.

The resulting histograms demonstrate an important and reassuring trend: the distribution of correlations for most features tends to follow a bell-shaped, approximately normal distribution, centered around the overall correlation observed for that feature in the entire dataset. For example, features such as *air temperature*, *clothing insulation*, and *metabolic rate* show per-user correlation distributions that are skewed in a direction that aligns with our general expectations:

- Air temperature: Most users exhibit a positive correlation between air temperature and thermal sensation, which reflects the intuitive relationship that warmer air generally increases the likelihood of a user reporting feeling warmer.
- Clothing insulation: The distribution of correlations is negatively skewed, suggesting that users who wear heavier clothing tend to feel cooler. This does not imply that heavier clothing causes cold discomfort but reflects behavioral adaptation—users dress more warmly when it is cold, and thus, heavier clothing is associated with colder perceived conditions.
- **Metabolic rate**: The distribution is positively skewed, consistent with the principle that higher activity levels lead to greater internal heat generation, increasing thermal sensation.

These patterns indicate that while there is individual variability, many of the observed relationships between features and thermal sensation are consistent across users. This supports the idea that a general model can capture the dominant thermophysiological responses, while also suggesting room for fine-tuning to capture individual nuances. This analysis contributes to validating our modeling approach: while personalization remains crucial for optimal comfort prediction, the dataset contains consistent and learnable patterns across users that justify the use of supervised learning techniques.



Figure 4.5: Distributions of per-user feature correlations with thermal sensation

4.2.6 Feature Engineering

The raw ASHRAE II dataset includes a wide variety of environmental, physiological, and contextual variables. For effective model training, several feature engineering strategies were employed to enhance model performance and ensure robustness.

Variable Selection: Features were selected based on their expected relevance to thermal sensation and their availability across the majority of records. We prioritized features that are commonly measurable in real-world environments or require minimal user interaction. The final set included:

• User Attributes: Age, gender (one-hot encoded), metabolic rate, clothing insulation.

- Indoor Environmental Features: Air temperature, relative humidity, air velocity.
- **Outdoor Environmental Features:** Outdoor air temperature and relative humidity.
- **Control States:** Fan and window state, to assess their role in low-cost occupant-controlled comfort adjustments.
- **Temporal Feature:** Day-of-year cosine transformation (detailed below).

Height and weight were explicitly excluded from the model due to their high correlation with metabolic rate and body mass index (BMI), leading to multicollinearity concerns. Moreover, their direct influence on thermal sensation is indirect and largely mediated through metabolic effects, which were already captured via the metabolic rate variable.

Handling Missing Values: Some features exhibited partial missingness. Notably, external humidity values were occasionally absent in the original dataset. In such cases, we substituted values from ISD (Integrated Surface Database) weather stations, specifically the rh_out_isd column, using spatial proximity to ensure representativeness. These substitutions led to improved model coverage and performance, without introducing significant noise.

Target Variable Preprocessing: The primary target was thermal sensation, recorded on the ASHRAE 7-point scale from -3 to +3. Although most values were discrete integers, a small number of floating-point labels were present due to interpolation or measurement rounding. Preliminary tests showed that rounding these to the nearest integer negatively affected classification performance, likely introducing inconsistency. Therefore, these non-integer records were excluded.

Temporal Encoding: To capture seasonal patterns in user comfort preferences, we encoded the day of the year (DOY) using a cosine transformation:

$$\operatorname{doy_cos} = \cos\left(2\pi \cdot \frac{\operatorname{DOY}}{365}\right)$$

This cyclical encoding assigns high values near winter, low values during summer, and intermediate values in spring and autumn. The cosine function was chosen over sine because it aligns its peaks and troughs with thermal relevance: cosine yields +1 during winter and -1 during summer, aligning well with typical seasonal comfort trends. In contrast, a sine transformation would peak in spring and reach its minimum in autumn, which does not directly reflect thermal discomfort extremes. This encoding helps the models generalize across seasonal boundaries and enhances their ability to infer comfort expectations from time alone.

For locations in the Southern Hemisphere, where the seasons are reversed, the transformation must be inverted to preserve thermal relevance. Specifically, the value of $\cos\left(2\pi \cdot \frac{\text{DOY}}{365}\right)$ should be negated (i.e., $-\cos(\ldots)$) to align winter with +1 and summer with -1. This adjustment ensures that the cyclical representation reflects true seasonal thermal effects globally, accommodating hemispheric differences in comfort dynamics.

Data Scaling: Each model type required different normalization techniques. For treebased models (Random Forest, XGBoost), no explicit scaling was necessary. For the MLP classifier, MinMaxScaler was applied to normalize features into the [0, 1] range, promoting numerical stability during backpropagation. For the CNN-LSTM model, **RobustScaler** was used. This scaler removes the median and scales data using the interquartile range, offering resilience to outliers. It is especially well-suited for physiological and comfort data where extremes can arise due to individual variability.

4.3 Baseline Models

To establish a comprehensive performance benchmark, we trained and evaluated four well-established classification models: Logistic Regression, Random Forest, XGBoost, and Multi-Layer Perceptron (MLP). These models serve both as initial predictors and as comparative baselines for more complex and personalized strategies explored later in this work.

Evaluation Metrics

The primary metric used for evaluating model performance is the **macro-averaged F1-score**, which is particularly suited for imbalanced multiclass classification problems. Thermal sensation labels in our dataset range from -3 (cold) to +3 (hot), with a heavy skew toward the neutral class (0), making accuracy a misleading measure.

The F1-score for a single class is defined as the harmonic mean of precision and recall:

$$F1 = 2 \cdot \frac{\text{Precision} \cdot \text{Recall}}{\text{Precision} + \text{Recall}}$$

In the macro-averaged version, the F1-score is computed independently for each class and averaged:

$$F1_{\text{macro}} = \frac{1}{C} \sum_{i=1}^{C} F1_i$$

where C is the total number of classes. This approach gives equal weight to each class, ensuring that rare classes (e.g., extreme thermal sensations like -3 or +3) are not ignored during optimization.

Model Selection Rationale

Logistic Regression was selected for its simplicity, interpretability, and widespread use in classification. It serves as a strong linear baseline that highlights limitations of non-personalized, feature-linear decision boundaries.

Random Forest was chosen due to its capability to handle non-linearities and its robustness to missing data and feature noise. It is known for its stability and interpretability through feature importance rankings.

XGBoost represents a more sophisticated ensemble method that incorporates gradient boosting and regularization, allowing it to uncover complex patterns that may elude Random Forests. Its superior performance on structured tabular data made it a prime candidate for this task.

MLP (Multi-Layer Perceptron) provides a non-linear, neural-based alternative capable of learning deep feature interactions. Although it requires more tuning and data preprocessing, it has the potential to outperform tree-based models when trained correctly.

Feature Relevance and Model Behavior

Across all models, indoor air temperature consistently emerged as the most influential predictor, aligning with thermal comfort literature. Other strong contributors included relative humidity, clothing insulation, and metabolic rate. Logistic Regression revealed these patterns linearly, whereas Random Forest and XGBoost exposed interactions and thresholds, such as drastic comfort changes beyond specific temperature ranges.

MLP's behavior highlighted its reliance on normalized and complete data. It was more sensitive to training settings (e.g., learning rate, number of hidden units) and required MinMaxScaler preprocessing, unlike the tree-based models that operated effectively on raw features.

Training and Data Challenges

All models were trained using stratified train-test splits to maintain class distribution, mitigating bias from the overrepresentation of the neutral class. Non-integer thermal sensation values, which negatively affected model calibration when rounded, were excluded rather than imputed.

Model training also confirmed that Random Forest and XGBoost were more resilient to small perturbations in data and hyperparameters, whereas MLP exhibited variability in convergence, necessitating multiple trials.

4.4 Experimental Results

This section presents the results of the core modeling experiments using four supervised learning algorithms: Logistic Regression (LR), Random Forest (RF), Multi-Layer Perceptron (MLP), and Extreme Gradient Boosting (XGBoost). Each model was trained on the same set of preprocessed and feature-engineered data, with performance evaluated using accuracy and macro-average F1 score across the 7-point thermal sensation scale.

These results serve a dual purpose. First, they demonstrate the predictive feasibility of data-driven models for personalized thermal comfort. Second, they validate experimental design choices by confirming consistency with prior findings in the literature. In particular, our results echo model hierarchies seen in related works such as [23] and [37], where ensemble methods like Random Forests and XGBoost outperform simpler baselines such as Logistic Regression.

4.4.1 Model Performance Overview

Each model was trained using a stratified 80/20 train-test split to preserve class distributions. Evaluation focused primarily on the **macro-average F1 score**, which is especially relevant for multiclass problems with class imbalance—such as the 7-point thermal comfort scale ranging from -3 (cold) to +3 (hot). Unlike accuracy, which may be inflated by the predominance of neutral votes (i.e., class 0), the macro F1 treats all classes equally by averaging the F1 scores for each class:

$$F1_{\text{macro}} = \frac{1}{N} \sum_{i=1}^{N} \frac{2 \cdot \text{precision}_i \cdot \text{recall}_i}{\text{precision}_i + \text{recall}_i}$$
(4.1)

Where N is the number of classes, and $\operatorname{precision}_i$ and recall_i are computed individually per class i.

The performance summary is illustrated in Table 4.1 and visually in Figure 4.6.

Table 4.1. Ferrormance Summary of Dasenne Models		
Model	Accuracy (%)	Macro Avg F1 Score
Logistic Regression	0.27	0.26
Random Forest	0.50	0.49
MLP	0.51	0.45
XGBoost	0.52	0.48

 Table 4.1: Performance Summary of Baseline Models



Figure 4.6: Cross-Model Comparison Bar Plot

4.4.2 Logistic Regression

The Logistic Regression model, while interpretable and computationally efficient, yielded the weakest overall performance. This is expected due to the inherent **non-linearity of thermal comfort perception**, which cannot be effectively captured by linear decision boundaries.

Despite attempts at regularization and balanced class weighting, the model consistently favored the dominant classes, particularly class 0 (neutral), leading to significant drops in precision and recall for less represented classes (e.g., +3 or -3).

4.4.3 Random Forest

The Random Forest classifier delivered the most stable and robust performance among all models. After an extensive hyperparameter search, the following configuration proved optimal:

- $n_{estimators} = 400$
- $max_depth = 15$ outperformed both shallow (10) and unrestricted depths
- min_samples_split = 5, min_samples_leaf = 1

Feature importance analysis (see Figure 4.7) showed that environmental features like **air temperature, clothing insulation**, and **metabolic rate** were the most informative, while **gender** contributed little predictive power. This agrees with recent findings advocating the importance of personal thermal exposure over demographic attributes. The overall macro F1 score remained modest, reflecting the difficulty of the classification task, but it surpassed other methods in both reliability and interpretability.



Figure 4.7: Random Forest Feature Importance

4.4.4 Multi-Layer Perceptron (MLP)

The MLP model was the most sensitive to hyperparameter configurations. The following settings led to the best results:

- Hidden layers: (300, 200, 20)
- Learning rate: 0.01 outperformed both 0.001 and 0.0001
- Regularization (alpha): 0.0001 struck a balance between underfitting and overfitting
- Activation: ReLU with Adam optimizer

Interestingly, models with too few (1 layer) or too many layers (4+) degraded sharply in performance, suggesting that moderate complexity is required to model the nonlinear patterns in thermal comfort perception without overfitting.

The model's performance approached that of XGBoost but exhibited higher variance across training runs.

4.4.5 XGBoost

XGBoost offered competitive performance and slightly exceeded MLP in macro F1. It was, however, more prone to overfitting, especially with high max_depth and low subsample values. Optimal hyperparameters included:

• $max_depth = 15$

- learning_rate = 0.1
- subsample = 0.7
- colsample_bytree = 1.0

Tuning learning_rate provided marginal gains in accuracy, but exacerbated variance and class imbalance sensitivity.

4.4.6 Confusion Matrix Analysis

The class-wise confusion matrices for each model (Figures 4.8–4.11) provide additional insights. Logistic Regression consistently confused extreme values (e.g., -3, +3), often collapsing them into nearby neutral categories. Random Forest and XGBoost captured broader variance but struggled with sparsely represented classes.



Figure 4.8: Confusion Matrix – Logistic Regression



Figure 4.9: Confusion Matrix – Random Forest


Confusion Matrix: MLP Classifier

Figure 4.10: Confusion Matrix – MLP



Figure 4.11: Confusion Matrix – XGBoost

These patterns reinforce the rationale for using macro F1 and highlight the continuing challenge of modeling rare but significant thermal comfort extremes.

4.4.7 Effect of Window Opening on Predicted Comfort

In addition to evaluating baseline model performance, a dedicated experiment was conducted to assess the role of operable windows as an actionable variable in personalized thermal comfort modeling. The window feature (binary: open = 0, closed = 1) was included alongside core environmental and user variables in the best-performing Random Forest classifier, which was retrained accordingly. The dataset used for this experiment incorporated the **outdoor relative humidity** values from ISD stations, replacing missing entries in the original dataset. This substitution was based on earlier preprocessing analysis indicating that ISD data improved robustness and coverage.

Although the window feature ranked lowest in feature importance in the Random Forest model (see Figure 4.12), this alone does not fully capture its impact. Many environmental variables exhibit strong multicollinearity, particularly with air temperature and humidity. Therefore, we designed a counterfactual experiment to test the window feature's causal influence on predicted comfort.

Counterfactual Analysis Design

Using the trained model, we generated two modified versions of the test set:

- One where all window values were set to open (0)
- One where all window values were set to closed (1)

Predictions were generated for both versions. The difference between predicted thermal comfort scores (i.e., $\Delta Comfort = Comfort_{open} - Comfort_{closed}$) was computed for each sample. The distribution of these differences is presented in Figure 4.13.

Results

Despite its low feature importance in the original model, the counterfactual analysis revealed that:

- 10% of samples showed a different predicted thermal sensation when toggling the window state.
- 6.17% of samples showed an increase in predicted thermal sensation due to the change.

The Δ Comfort distribution was centered around zero, but with a long tail suggesting localized effects in specific environmental or seasonal contexts. These results imply that window operations may have non-negligible influence in comfort prediction, particularly in certain boundary conditions. However, the limited dataset size (only 1,411 samples after filtering) and strong correlation with air temperature likely constrained its standalone predictive power.

To further assess the effect of window state on thermal comfort, we evaluated whether switching the window setting brought the predicted comfort value *closer to neutral* (i.e., closer to zero on the thermal sensation scale). This approach recognizes that movement toward neutrality, not just raw change, is a more meaningful indicator of improved comfort.

In this refined analysis:

- 5.10% of predictions moved *closer* to the neutral value (0) when the window was set to open.
- **3.83%** of predictions moved *further away* from neutral.

The distribution of these proximity deltas is illustrated in Figure 4.14. It shows that while most predictions remained unaffected, a measurable subset experienced improved proximity to comfort when the window was opened—supporting the view that such control variables can yield context-sensitive gains in occupant satisfaction.



Figure 4.12: Feature importance from Random Forest model including the window variable



Figure 4.13: Change in predicted comfort (Δ Comfort) under window state alteration (counterfactual test)



Figure 4.14: Change in distance to neutral comfort score ($|Comfort_{closed}| - |Comfort_{open}|$). Positive values indicate improved proximity to thermal neutrality when window is open.

Relevance to Smart Environments

In the broader context of smart environments, the inclusion and evaluation of controllable features like window state is of particular practical relevance. Unlike static user attributes (such as age or gender) or immutable environmental variables (like regional climate), operable elements such as windows represent **direct points of interaction between users and their environment**. This positions them as key components in occupant-centered adaptive systems.

The findings of the counterfactual experiment—where 9% of predictions changed and 6% indicated improved comfort upon altering window state—highlight the **latent potential** of passive, low-cost interventions. Even though window status received low importance in traditional feature attribution, its situational influence becomes apparent through scenario-based testing.

Such an approach aligns well with the goals of intelligent building management systems, which aim not only to optimize energy efficiency but also to dynamically adapt to user needs with minimal manual input. From a systems design perspective, integrating actuator-level feedback into control logic could lead to more **personalized and explain-able comfort recommendations**, contributing to the growing field of human-centered automation. Thus, the experiment underscores that **actuator features—while subtle in global models—may still yield significant value in personalized and context-aware smart control frameworks**.

4.4.8 Geographical Subsetting and Model Performance

To investigate the influence of geographical and climatic contexts on model performance, the Random Forest classifier was evaluated across various subsets of the dataset, stratified by country, city, and climate zone. The performance metric utilized was the macroaveraged F1-score, which provides a balanced measure of predictive accuracy across all classes, regardless of class imbalance.

Performance Variability Across Regions

The evaluation revealed notable variability in model performance across different geographical subsets:

- High F1 Macro Scores:
 - Pakistan (Country): Achieved an F1 macro score of 0.610 with 4,068 samples.
 - Desert (Hot Arid) Climate: Recorded an F1 macro score of 0.522 with 2,011 samples.

These high scores suggest a more homogeneous population and environmental setting, where the combination of sufficient sample size and low intra-group variance likely contributed to the model's ability to generalize effectively.

- Low F1 Macro Scores:
 - Bangkok (City): Recorded an F1 macro score of 0.305 with 1,125 samples.
 - Brisbane (City): Achieved an F1 macro score of 0.298 with 1,104 samples.
 - *Humid Subtropical Climate:* Recorded an F1 macro score of 0.298 with 1,104 samples.
 - *Tropical Savanna Climate:* Achieved an F1 macro score of 0.263 with 3,379 samples.
 - Townsville (City): Recorded an F1 macro score of 0.256 with 1,211 samples.
 - Australia (Country): Achieved an F1 macro score of 0.255 with 5,045 samples.

Despite moderate to high data volumes, these regions exhibited lower model performance. This may be attributed to greater user diversity in comfort perception, environmental variability, particularly in mixed or transitional climates, and insufficient sample size to capture complex patterns across diverse subgroups.

Implications and Recommendations

These findings underscore that geographical or climatic labels alone do not guarantee homogeneity. For instance, Australia's large territory spans diverse climatic zones and demographic profiles, which may dilute model accuracy. Similarly, the "tropical savanna" label aggregates multiple regions and cultures, reducing internal consistency.

To enhance model performance, the following strategies are recommended:

- Stratify by More Granular Subgroups: Consider user profiles within climates to capture more homogeneous subpopulations.
- Explore Clustering Within Large Datasets: Utilize clustering algorithms to identify homogeneous subpopulations within diverse datasets.
- Analyze Intra-Group Variance of Key Features: Examine the variance of features such as thermal sensation, temperature, and humidity within groups to assess homogeneity.
- Consider Dimensionality Reduction or Unsupervised Learning: Apply techniques like t-Distributed Stochastic Neighbor Embedding (t-SNE) to detect hidden group structures.

Conclusion

The geographical subsetting analysis reveals that model performance in thermal comfort prediction is significantly influenced by regional and climatic factors. High performance in regions like Pakistan and desert climates suggests that homogeneity in environmental conditions and user profiles enhances model generalizability. Conversely, lower performance in regions with greater diversity and environmental variability highlights the need for more nuanced modeling approaches that account for intra-group differences. Future work should focus on developing adaptive models that can dynamically adjust to the specific characteristics of different geographical and climatic contexts.

4.5 Transfer Learning for Climate-Adaptive Modeling

This section explores the potential of transfer learning (TL) to address challenges of generalizability and data scarcity in thermal comfort modeling. While classical models rely on direct feature-target mappings within a single domain, TL offers the ability to reuse knowledge from well-represented domains (source climates) to enhance model performance in underrepresented or more complex environments (target climates). This is particularly relevant in the thermal comfort domain, where subjective responses are influenced by a combination of personal, environmental, and contextual variables that may not be consistently available across locations.

We follow the architecture proposed in [23], which combines convolutional and recurrent neural layers to learn both spatial and temporal representations of thermal comfort data. The core idea is to train the model on a large and relatively homogeneous climate zone and fine-tune it to a different target climate.

4.5.1 Model Design and Implementation

To investigate the potential of transfer learning in personalized thermal comfort modeling across climates, we implemented a deep learning model based on the CNN-LSTM architecture described by Somu et al.[23]. This architecture is designed to handle spatiotemporal input and is particularly suited to capture the sequential patterns present in occupant comfort data collected over time.

The model architecture begins with a **1D** convolutional layer that processes each time step in the sequence of environmental and physiological measurements. This layer uses **128 filters**, each with a **kernel size of 5**, and 'same' padding, maintaining the temporal length of the input. A **spatial dropout** of **0.1** is applied post-convolution to reduce overfitting.

The convolutional output, shaped as a sequence of 128-dimensional feature vectors, is passed through **two stacked LSTM layers**, each comprising **256 hidden units**. The first LSTM layer includes a **recurrent dropout of 0.1** to further regularize the learning process, while the second LSTM layer operates without dropout. The LSTM layers model temporal dependencies across the sequence, allowing the model to capture evolving patterns in user comfort states.

The output from the final timestep of the LSTM is fed into a **two-layer fully connected neural network**, with **64** and **16** units respectively, both using **ReLU activation**. A

final softmax output layer with 7 neurons corresponds to the full ASHRAE 7-point thermal sensation scale, ranging from cold (-3) to hot (+3).

Training was conducted using the Adam optimizer with a learning rate of 0.001, and categorical cross-entropy loss. The model was trained over 30 epochs on a total of 8,349 sequence samples, derived from users with sufficient temporal data. The sequence length was fixed at 5, meaning each training example consisted of 5 consecutive records per subject.

Data Preparation: Before training, the data were **grouped by subject** and **sorted chronologically**. Time differences between consecutive measurements were computed, revealing that most samples within a user occurred over short periods (typically within a few days). This temporal compactness is advantageous for sequence-based models like LSTMs, which are designed to capture local temporal dynamics. Features were scaled using **RobustScaler**, ensuring stability in the presence of outliers.

Justification for Full 7-Class Prediction: Unlike the referenced study, which merged datasets and downsampled the output space to a 5-point scale, we maintained the original 7-class scale. This choice was based on earlier findings where even minor ambiguities—such as non-integer thermal sensation values—reduced model performance. Downsampling would likely merge adjacent but semantically distinct classes (e.g., +2 "warm" with +1 "slightly warm"), reducing the resolution of predictions. Maintaining the full scale allowed the model to learn finer-grained distinctions, albeit at the cost of increased classification difficulty.

4.5.2 Training and Evaluation on the Source Domain

To construct a robust foundation for subsequent transfer learning across climate zones, we first trained the CNN-LSTM architecture on the full ASHRAE II dataset. This phase aimed to enable the model to learn generalized spatiotemporal patterns in thermal comfort behavior based on real-world, longitudinal data across diverse users, buildings, and environmental contexts.

The training dataset was built using a sliding window mechanism to exploit the sequential nature of occupant comfort data. Specifically, individual occupant records were grouped by subject_id and temporally ordered using their timestamp fields. The dataset's time resolution, generally at a daily granularity, supported the creation of meaningful short-term temporal sequences. For each subject with sufficient data, sequences of length five consecutive time steps were generated, representing approximately a working week's worth of environmental and physiological observations. The thermal sensation label associated with the final sample in each sequence was used as the prediction target. This aligns with the formulation in the CNN-LSTM transfer learning study [23], which demonstrated the effectiveness of leveraging both convolutional and recurrent components for capturing localized temporal trends and long-term dependencies, respectively.

Following this preprocessing pipeline, the final dataset consisted of 8,349 training sam-

ples and 2,088 validation samples, totaling 10,437 sequences. These sequences encompassed 9 features per timestep: age, metabolic rate, clothing insulation, air temperature, relative humidity, air velocity, outdoor temperature, outdoor relative humidity, and gender (as a binary indicator). These variables were selected for their theoretical and empirical relevance to thermal comfort, as evidenced in the ISO 7730 and ASHRAE 55 standards and in recent machine learning literature. Each 5-timestep sequence was represented as a 5×9 matrix and appropriately reshaped to match the model's convolutional input expectations.

To ensure scale consistency and robustness to outliers, we applied the **RobustScaler** from scikit-learn to all continuous input features. This transformation subtracts the median and scales the data by the interquartile range, mitigating the influence of extreme values or skewed distributions—an important property given the variability in real-world sensor data.

The model was trained using the Adam optimizer with a learning rate of 0.001, a batch size of 128, and for 30 epochs. A fixed 80/20 train-validation split was used, with stratification based on the 7-point thermal sensation scale to preserve class distribution across both subsets. No oversampling or augmentation was used, which potentially contributed to the underrepresentation of extreme classes during training.

Early training epochs showed promising convergence, with training loss decreasing and accuracy increasing steadily. However, beginning around epoch 10, a growing divergence between training and validation loss was observed. This is illustrated in Figures 4.15 and 4.16. While the training loss continued to decrease, the validation loss began to increase, suggesting the onset of overfitting. Similarly, training accuracy climbed steadily while validation accuracy plateaued.

Final performance metrics reflected this dynamic: the model achieved an **accuracy of 51.3%** and a **macro-averaged F1 score of 0.385** on the validation set. While the accuracy is consistent with results from simpler models (e.g., Random Forests), the relatively lower macro F1 score suggests a difficulty in achieving balanced performance across all classes—especially the edge cases of -3 and +3 thermal sensations.

One important distinction between our implementation and the referenced study is the preservation of the original 7-point thermal sensation scale. In contrast, [23] consolidated the classes into a 5-point scale, reducing complexity and label imbalance. Our decision favored interpretability and fidelity to the original dataset but may have adversely affected performance by increasing inter-class ambiguity.

In summary, while the source-domain model effectively captured useful general patterns, its performance was limited by class imbalance and label granularity. Nevertheless, the latent features learned in its early layers are expected to support more efficient adaptation to new environments in the target domain—a hypothesis we evaluate in the next section.



Figure 4.15: Training vs. Validation Loss across epochs for CNN-LSTM on ASHRAE dataset

4.5.3 Transfer Learning to Target Climate

To investigate the efficacy of transfer learning in climate-specific thermal comfort prediction, the pretrained CNN-LSTM model was adapted to a specific climatic region using



Figure 4.16: Training vs. Validation Accuracy across epochs for CNN-LSTM on ASHRAE dataset

domain-focused fine-tuning. The tropical savanna climate was selected as the target due to its high representation within the ASHRAE II dataset. After preprocessing and filtering, the subset contained 1,910 samples—sufficient to explore transfer learning while still reflecting real-world constraints of limited, non-uniform user data across climate zones.

This transfer learning experiment followed the core idea of reusing knowledge gained in one domain (source climate) to improve predictions in another (target climate). Drawing from the methodology outlined in [23], the adaptation procedure involved freezing the deeper, fully connected layers of the trained model—those believed to encode high-level abstract knowledge about thermal sensation—and retraining only the earlier convolutional and recurrent (LSTM) layers. This setup assumes that lower-level features such as the temporal evolution of environmental signals may vary more across climates, while the mapping to thermal sensation categories remains relatively consistent.

The model was fine-tuned for 10 epochs using the standard training loop, loss function, and learning rate previously applied. Importantly, the target climate data was not excluded from the original source model training set, meaning this approach does not fully represent a strict domain-adaptation setup. However, it does reflect a realistic use case where all available data is leveraged for general training, and model specialization is subsequently pursued via targeted refinement.

Despite the seemingly adequate volume of data, the model began to overfit rapidly—training accuracy improved steadily, while validation accuracy plateaued early and validation loss began diverging from training loss after only three epochs. This behavior is visually supported by the training and validation loss and accuracy curves plotted during training (see Figures 4.17 and 4.18).

Multiple hypotheses can explain this outcome:

- Class Imbalance and Oversampling Effects: The thermal sensation labels remain highly imbalanced in the target set. Although stratified sampling was used, no synthetic oversampling techniques were applied. The difficulty of distinguishing neighboring classes on the 7-point ASHRAE scale may have exacerbated learning instability.
- Insufficient Intra-Group Homogeneity: Although labeled under the same climate category, the tropical savanna group may contain geographically and demographically diverse populations, weakening the model's ability to detect consistent spatio-temporal patterns.
- Model Complexity vs. Data Granularity: The CNN-LSTM architecture, while powerful, may be overly complex relative to the granularity and variability of thermal comfort data. Unlike image or speech datasets, thermal comfort features are limited in number and heavily influenced by subjective, untracked psychological or contextual factors.

Nonetheless, this transfer learning experiment remains valuable. It underscores the challenges of applying deep learning architectures in thermal comfort prediction, especially when targeting real-world deployment across heterogeneous climatic settings. Future work should consider pretraining on strictly separated source domains and adopting simplified target-specific architectures or regularization techniques to prevent early overfitting.



Figure 4.17: Training vs. Validation Loss – Transfer Learning on Tropical Savanna Climate



Figure 4.18: Training vs. Validation Accuracy – Transfer Learning on Tropical Savanna Climate

4.5.4 Discussion and Implications

The transfer learning experiment revealed both potential and limitations of applying deep spatio-temporal models in the context of climate-adaptive thermal comfort modeling. While the CNN-LSTM model achieved a comparable performance to the best classical models when trained on the full dataset, the transfer to a new target domain—specifically the tropical savanna climate—resulted in immediate overfitting and significantly degraded performance.

This outcome underscores the sensitivity of complex neural architectures to data volume and class balance. Despite having over 1900 samples, the tropical savanna subset likely contained considerable intra-class variation. The use of the full 7-point ASHRAE thermal sensation scale, rather than downsampling to 5 classes as in the reference study[23], further increased the classification difficulty. As a result, the model may have struggled to learn meaningful distinctions between neighboring classes, especially in the presence of user heterogeneity and environmental noise.

The rapid divergence between training and validation curves highlights a critical challenge in transfer learning for comfort modeling: the pre-trained source domain knowledge can quickly become overly specialized, even when the low-level layers are adapted to the target domain. Although the freezing strategy followed the recommended structure from the reference paper, the outcome suggests that further experimentation with freezing strategies and hybrid training regimes might be necessary to improve generalization.

Practically, this experiment highlights the trade-off between model complexity and data adequacy. It suggests that deep learning models with transfer learning may only outperform simpler models when target domain data are sufficiently abundant and homogenous. For personalized applications in smart environments, this reinforces the need for scalable model architectures that can adjust to varying levels of user and environmental data, or alternatively for hybrid systems that blend neural representations with rule-based personalization.

Chapter 5

Discussion

5.1 Summary of Modeling Results

This section consolidates the performance outcomes from the diverse machine learning models developed throughout the study. These results illustrate the challenges, strengths, and limitations associated with modeling thermal comfort using the ASHRAE II dataset — a complex, real-world dataset capturing diverse environmental, demographic, and subjective factors.

The **Random Forest classifier** consistently emerged as the most effective among the baseline models. With a macro-averaged F1 score ranging between 0.47 and 0.51 across different experiments, it demonstrated solid predictive ability even in the face of data imbalance and class subjectivity. Random Forests were particularly well-suited due to their ensemble nature, inherent feature selection capability, and robustness to multicollinearity. Their ability to capture nonlinear relationships, while also remaining interpretable through feature importance scores, made them ideal for extracting actionable insights.

The feature importance analysis from Random Forests provided clarity on the driving factors of thermal sensation. As expected, *air temperature*, *clothing insulation*, and *metabolic rate* were among the top predictors. Interestingly, demographic variables such as *gender* showed minimal impact, while *age* displayed moderate influence. The inclusion of a *seasonal cosine variable* (representing the day of year) also proved valuable, capturing cyclic temporal patterns linked to comfort shifts across seasons.

XGBoost, a gradient-boosted decision tree model, yielded results comparable to the Random Forest but with higher sensitivity to hyperparameter choices. When tuned appropriately (e.g., max_depth=15, subsample=0.7), XGBoost achieved slightly higher accuracy, though it tended to overfit more readily. This suggests that while boosting techniques offer strong predictive potential, their use in thermal comfort prediction requires careful regularization, especially in datasets with noisy, overlapping classes and unbalanced label distributions.

Multilayer Perceptrons (MLPs) offered a neural alternative, with meaningful gains in performance once suitable architectures and learning parameters were established. A three-layer architecture with nodes arranged as (300, 200, 20) and a learning rate of 0.01 yielded the best results. However, MLPs demonstrated higher variance in training stability and required rigorous preprocessing and scaling. Moreover, unlike Random Forests, they lacked explainability, making their practical deployment in building systems more challenging.

Logistic Regression, as expected, underperformed significantly. It failed to model the complex, nonlinear nature of the input space, yielding a macro F1 score well below 0.35. This outcome reinforces the fact that thermal comfort data are not linearly separable and that linear models are insufficient to capture nuanced interactions between environmental conditions, personal attributes, and subjective responses.

The evaluation centered on the **macro-averaged F1 score**, a metric particularly suitable for this study due to the **strong class imbalance**. The dataset showed a clear dominance of the "neutral" class (value 0), which could mislead accuracy-based assessments. Macro F1 treats each class equally, computing the unweighted mean of the F1 scores across all classes, thereby providing a more balanced and fair evaluation of model performance.

Collectively, these results validate the use of machine learning in modeling thermal comfort. However, they also underscore the **inherent limitations** of the task: the subjectivity of comfort, the fine-grained nature of the 7-class scale, and the limited ability of environmental and demographic data to capture psychological or behavioral influences. These challenges motivate the exploration of personalized models and transfer learning approaches in subsequent sections.

5.2 Limitations of Model Performance

Despite the adoption of a diverse set of machine learning approaches, from classical classifiers to deep learning-based architectures, a range of practical and theoretical limitations constrained the performance of the models. These limitations stem from the complexity of the problem domain, the nature of the data collected, and methodological trade-offs made during modeling.

Subjectivity of the Target Variable

Thermal comfort is not only a physiological phenomenon but also a deeply subjective experience. The ASHRAE 7-point thermal sensation scale, used as the target variable, captures individual perceptions ranging from "cold" to "hot." These ratings are inherently personal and influenced by behavioral, psychological, and cultural factors not captured in the dataset. Consequently, even under identical environmental and demographic conditions, two individuals may report differing comfort levels. This inter-subject variability introduces substantial noise into the learning task and limits the attainable predictive accuracy, especially for generalized models.

Fine-Grained and Overlapping Class Structure

The 7-point scale, while useful for granular assessment, introduces significant modeling

complexity. Adjacent categories such as "slightly cool," "neutral," and "slightly warm" often exhibit overlapping feature distributions, making them difficult to distinguish even for sophisticated models. Misclassification between these classes is not only more probable but also harder to penalize meaningfully, as the thermal comfort boundaries are not sharp. Furthermore, rounding or downsampling (as avoided in this study) could simplify the classification task but would compromise the interpretability and fidelity of predictions.

Class Imbalance and Data Sparsity

The dataset reflects natural reporting tendencies, with a skew toward 'neutral' responses, while extreme sensations like 'hot' (+3) or 'cold' (-3) are rarely observed. This imbalance impairs the ability of models to learn discriminative features for minority classes. Although strategies such as class weighting and macro-average metrics (F1 score) were applied to mitigate this, the fundamental issue of insufficient class representation remains a bottleneck. In addition, users with few records were excluded from many experiments, further reducing the sample space available for underrepresented groups.

Unobserved and Unmeasurable Influences

Thermal comfort is influenced by many variables not present in the ASHRAE dataset. These include physical exertion prior to measurement, emotional state, acclimatization, hydration level, and recent exposure to different environments. While the dataset does record basic physiological indicators such as clothing insulation and metabolic rate, it lacks real-time biometric feedback or contextual behavioral cues. These missing variables limit the model's ability to capture the full causal structure behind thermal perception.

Temporal Inconsistencies in Sequential Modeling

In the sequence-based modeling (CNN-LSTM), data were grouped per user and ordered by timestamp. Although most records occurred in relatively short intervals (e.g., within the same day), the sampling frequency varied significantly across users and studies. This irregularity challenges the core assumptions of LSTM-based models, which typically benefit from regularly spaced sequences. Furthermore, fixed-length windows may include redundant data or omit important transitional states, depending on the user's record density.

Model Complexity vs. Generalizability

While deeper architectures such as MLP and CNN-LSTM provided greater modeling capacity, they also risk overfitting—especially when the number of samples per user or region was limited. This tension between model complexity and dataset size became particularly evident in transfer learning experiments, where initial overfitting appeared within just a few epochs. Even advanced techniques such as layer freezing and temporal augmentation could not fully mitigate this, highlighting the sensitivity of personalized models to training volume and heterogeneity.

Together, these limitations clarify the need for targeted strategies that go beyond generalpurpose modeling. Approaches such as personalized modeling, context-aware prediction, and integration of wearable data are promising directions that could address these constraints in future iterations of thermal comfort systems.

5.3 Practical Implications for Smart Environments

The experimental findings of this study provide strong indications that machine learningbased thermal comfort modeling holds substantial promise for practical deployment in smart environments. These insights extend beyond model accuracy, revealing how predictive personalization and data-driven inference can be embedded into everyday spaces to support comfort, efficiency, and autonomy.

Deployability of Small-Scale Models. The success of compact machine learning models—especially the Random Forest and MLP classifiers—demonstrates the feasibility of implementing thermal comfort prediction in real-world applications. These models performed well with modest computational requirements, making them suitable for deployment in embedded systems, microcontrollers, or low-power IoT devices. This is particularly relevant for building automation systems where computational resources may be constrained but responsiveness and autonomy are critical.

Importantly, the models achieved reliable performance without requiring high-resolution biometric data or deep personalization. This supports a scalable deployment strategy where a basic profile (age, gender) combined with ambient sensor readings can already enable effective comfort prediction, reducing user burden and preserving privacy.

Environmental Feature Relevance for Sensor-Driven Systems. Feature importance analysis confirmed that ambient environmental factors such as indoor air temperature, relative humidity, and air velocity were the most influential predictors. Outdoor conditions, particularly temperature and humidity, also contributed valuable information. These findings validate a sensor-driven approach to comfort estimation, where relatively low-cost and commercially available environmental sensors can provide all the necessary inputs to make personalized predictions.

The inclusion of day-of-year cosine as a seasonal proxy also proved beneficial, suggesting that temporal features can improve comfort estimation without directly collecting timeintensive data such as occupancy schedules or adaptive thermal histories.

Operable Controls as Active Comfort Agents. The counterfactual analysis of the window feature demonstrated that actionable environmental features may exert influence beyond their statistical importance in standard feature rankings. While window state had low feature importance, manipulating it altered the predicted thermal comfort in 9% of cases, with 6% showing improvement. This suggests that smart systems should not disregard low-ranked features if they correspond to controllable interventions. Instead, they could be incorporated into user feedback loops or automated control logics, especially in naturally ventilated or hybrid buildings.

Comfort-Aware Energy Management. Beyond occupant satisfaction, thermal comfort prediction enables more intelligent energy use. By anticipating comfort violations before they occur, systems can proactively modulate HVAC operation. For example, minor deviations from neutral comfort might be tolerated to conserve energy, while predicted discomfort can trigger preemptive corrections. This comfort-aware regulation allows building managers or control algorithms to balance energy efficiency with user experience, moving beyond simple temperature thresholds.

Integration into Broader IoT Ecosystems. Given their lightweight computational demands and compatibility with standard sensor modalities, the tested models can be embedded within edge computing architectures. Comfort models could run on local gateways or room-level controllers, ingesting real-time sensor data and providing actionable outputs without cloud reliance. This architecture supports privacy-respecting, low-latency, and offline-capable smart environments, which are essential for residential, educational, or healthcare deployments where network access or data sensitivity may be concerns.

5.4 Personalization Potential and User-Centric Modeling

One of the central challenges in thermal comfort modeling is the inherent subjectivity and variability of human perception. Even under identical environmental conditions, different individuals can experience vastly different sensations of thermal comfort. This phenomenon was consistently observed throughout our experiments, with the same combination of temperature, humidity, and air velocity eliciting opposing comfort labels across subjects. Such discrepancies underscore the limitations of purely generalized models and highlight the promise of personalization in this domain.

Personalized thermal comfort models offer the potential to significantly improve prediction accuracy and occupant satisfaction. By tailoring predictions to an individual's historical comfort responses, physical characteristics, or behavioral patterns, these models can adapt to nuanced preferences that generalized models overlook. In real-world deployments, this can enable HVAC systems to deliver more targeted and energy-efficient conditioning, enhancing user well-being while minimizing unnecessary consumption.

However, realizing personalization presents key challenges. Chief among them is the data requirement: effective user-specific modeling demands a sufficient number of labeled observations per user. While our dataset showed promise, with over 50 subjects having more than 100 labeled records, such granularity is rare in typical deployments. Moreover, relying on explicit user feedback is impractical in daily life. Users are unlikely to tolerate frequent interruptions for feedback queries, making continuous and scalable feedback collection difficult.

Addressing this issue requires rethinking how user data is gathered. One potential approach is to **limit explicit feedback to a cold-start phase**, where a few labeled data points help bootstrap a personalized model. Beyond this, systems should shift to **passive or indirect data collection**, utilizing sensors and wearables. For instance, advanced smartwatches capable of measuring skin temperature, heart rate variability, or perspiration could offer valuable proxies for comfort state, enabling adaptation without active user input.

Another direction is to group users based on comfort profiles through unsupervised learn-

ing or clustering. Users with similar preferences or physiological traits can be aggregated, allowing for hybrid models that blend personalization with scalability. This strategy was suggested by our geographic and subgroup analyses, where performance improved in more homogeneous populations.

Ultimately, personalization must be implemented in a user-centric and privacy-aware manner. Data minimization, secure processing, and opt-in consent are critical to ensuring trust and acceptance. Nonetheless, the potential benefits—greater satisfaction, energy savings, and user empowerment—justify the exploration of this promising frontier.

Chapter 6

Future Work

6.1 Enhancing Comfort Prediction through Richer Data Sources

Thermal comfort is a multifaceted and subjective experience influenced by environmental conditions, physiological states, and individual behaviors. While models based on basic environmental parameters and user-reported features can yield acceptable performance, their predictive accuracy is inherently limited without deeper personalization and contextual awareness.

Recent advancements in wearable technology have enabled the continuous, non-invasive monitoring of physiological signals such as skin temperature, heart rate variability (HRV), and electrodermal activity. These physiological indicators have been shown to correlate significantly with thermal comfort perceptions. For instance, Lee and Chun (2021) developed a thermal comfort prediction model using physiological signals from wearable devices, achieving an accuracy of 80% with only physiological data [6]. Similarly, Nkurikiyeyezu et al. (2020) demonstrated that HRV could reliably predict thermal comfort states with up to 93.7% accuracy [7].

Incorporating additional environmental parameters can further enhance comfort prediction models. CO_2 concentration, for example, serves as a proxy for indoor air quality and ventilation effectiveness. Elevated CO_2 levels have been associated with decreased cognitive performance and increased discomfort [38]. Monitoring CO_2 concentrations alongside temperature and humidity can provide a more comprehensive understanding of indoor environmental quality.

Temporal patterns and occupant behaviors also play a crucial role in thermal comfort. Integrating data on occupancy schedules, activity levels, and diurnal cycles can enable models to anticipate comfort needs proactively. Long-term monitoring studies have shown that personal comfort models benefit from incorporating such temporal and behavioral data, leading to improved prediction accuracy [10].

The integration of these diverse data sources necessitates the use of sophisticated mod-

eling techniques capable of handling multimodal inputs. Machine learning approaches, such as ensemble methods and deep learning architectures, can effectively process and learn from complex datasets. Future research should focus on developing models that can seamlessly integrate physiological, environmental, and behavioral data to provide personalized and context-aware thermal comfort predictions.

In conclusion, enhancing thermal comfort prediction models through the integration of richer data sources holds significant promise for the development of intelligent, occupantcentric building systems. By leveraging physiological signals, environmental parameters, and behavioral patterns, these models can deliver more accurate and personalized comfort assessments, ultimately contributing to improved occupant well-being and energy efficiency.

6.2 Transfer Learning and Domain Adaptation

6.2.1 Motivation for Transfer Learning

Traditional thermal comfort models often require extensive labeled data, which may not be available for all building types or climates. Transfer learning allows models trained in one domain (e.g., a specific building or climate) to be adapted to another, reducing the need for large datasets in the target domain. This approach is particularly beneficial when deploying models in new environments with limited data availability.

6.2.2 Domain Adaptation Techniques

Several domain adaptation techniques have been explored to enhance the applicability of thermal comfort models:

- Unsupervised Domain Adaptation: Aligning feature distributions between source and target domains without labeled data in the target domain. Yang et al. [12] demonstrated the effectiveness of unsupervised domain adaptation techniques, such as Correlation Alignment (CORAL) and Dynamic Adversarial Adaptation Network (DAAN), in improving personalized thermal comfort predictions.
- Fine-Tuning: Adjusting pre-trained models using a small amount of labeled data from the target domain to improve performance. Gao et al. [11] applied fine-tuning in their transfer learning-based multilayer perceptron model for accurate thermal comfort prediction across multiple cities.

6.2.3 Applications in Thermal Comfort

Transfer learning has been applied in various thermal comfort scenarios:

• **Cross-Building Adaptation**: Applying models trained in office buildings to residential settings. Yang et al. [12] utilized unsupervised domain adaptation to transfer models between different building types. • **Cross-Climate Adaptation**: Adapting models from temperate to tropical climates, accounting for differences in occupant behavior and building design. Gao et al. [11] explored transfer learning for thermal comfort prediction in multiple cities within the same climate zone.

6.2.4 Future Directions

Future research should focus on:

- Investigating the effectiveness of different transfer learning strategies in various contexts.
- Developing standardized protocols for model adaptation across diverse environments.
- Exploring hybrid approaches that combine transfer learning with other machine learning techniques to enhance model robustness and accuracy.

6.3 Clustering and Segmentation for Personalized Modeling

6.3.1 Importance of User Segmentation

Thermal comfort is a subjective experience influenced by various factors, including age, gender, metabolic rate, clothing insulation, and individual preferences. Recognizing these differences is crucial for developing personalized thermal comfort models. Traditional models often assume a homogeneous occupant population, leading to generalized solutions that may not cater to individual needs. By segmenting users based on shared characteristics or behaviors, it's possible to tailor environmental controls more effectively, enhancing occupant satisfaction and energy efficiency.

6.3.2 Clustering Techniques

Clustering is an unsupervised machine learning approach that groups data points based on similarity. In the context of thermal comfort, clustering can identify groups of occupants with similar comfort preferences or behaviors. Common clustering techniques include:

- **K-Means Clustering**: This algorithm partitions data into k clusters by minimizing the variance within each cluster. It's widely used due to its simplicity and efficiency. For instance, researchers have applied K-Means to cluster occupancy profiles and energy demand patterns, aiding in the development of representative comfort models [39].
- **Hierarchical Clustering**: This method builds a tree-like structure (dendrogram) to represent nested groupings of data points. It's beneficial for understanding the relationships between different clusters and determining the optimal number of clusters [40].

- **Fuzzy Clustering**: Unlike hard clustering methods, fuzzy clustering allows data points to belong to multiple clusters with varying degrees of membership. This is particularly useful when occupant preferences overlap or are not distinctly separable.
- Gaussian Mixture Models (GMMs): GMMs assume that data points are generated from a mixture of several Gaussian distributions. They are effective in modeling data with subpopulations and have been used to classify energy and thermal comfort profiles in office buildings [40].

6.3.3 Implementation in Smart Environments

Implementing clustering techniques in smart environments involves several steps:

- 1. **Data Collection**: Gather data on environmental conditions (temperature, humidity), occupant characteristics (age, gender), and behaviors (occupancy patterns, clothing insulation).
- 2. Feature Selection: Identify relevant features that influence thermal comfort for clustering analysis.
- 3. Clustering Analysis: Apply appropriate clustering algorithms to segment occupants or spaces based on the selected features.
- 4. **Model Development**: Develop personalized thermal comfort models for each cluster, considering the specific preferences and behaviors of the group.
- 5. **HVAC Control Integration**: Tailor HVAC settings to meet the comfort requirements of each cluster, optimizing energy use and occupant satisfaction.

For example, a study utilized clustering to segment occupants based on their self-assessed thermal preferences (warmer, neutral, colder) and developed personalized comfort models for each group, resulting in improved prediction accuracy compared to generic models [5].

6.3.4 Challenges and Considerations

While clustering offers promising avenues for personalization, several challenges must be addressed:

- **Determining Optimal Cluster Numbers**: Selecting the appropriate number of clusters is critical. Over-segmentation can lead to overly complex models, while under-segmentation may overlook significant differences among occupants.
- **Dynamic Preferences**: Occupant preferences can change over time due to factors like acclimatization or seasonal variations. Models must be adaptable to such changes.
- **Data Privacy**: Collecting and analyzing personal data raises privacy concerns. Ensuring data anonymization and compliance with privacy regulations is essential.

• System Complexity and Cost: Implementing personalized models increases system complexity and may incur higher costs. Balancing personalization benefits with implementation feasibility is necessary.

6.4 Model Simplification and Edge Deployment

Deploying thermal comfort models on edge devices, such as smart thermostats and HVAC controllers, necessitates models that are both lightweight and efficient. These devices often have limited computational resources, memory, and power availability. Therefore, developing simplified models that can operate effectively within these constraints is crucial for real-time responsiveness and reliability.

Several techniques have been explored to reduce the complexity of machine learning models without significantly compromising their performance:

- **Pruning**: This technique involves removing less significant weights or neurons from a neural network, effectively reducing its size and computational requirements. Pruning can be structured or unstructured, with structured pruning often leading to more efficient implementations on hardware [31].
- Quantization: Quantization reduces the precision of the model's parameters, typically converting 32-bit floating-point numbers to 8-bit integers. This reduction leads to smaller model sizes and faster computations, making it suitable for deployment on resource-constrained devices [32].

Implementing simplified models directly on edge devices offers several advantages, including reduced latency, improved privacy, and decreased reliance on cloud connectivity. Strategies for effective edge deployment include:

- **Integration with HVAC Controllers**: Embedding models within HVAC systems allows for real-time adjustments based on occupant comfort preferences, leading to enhanced energy efficiency and user satisfaction.
- Utilization of Smart Thermostats: Smart thermostats equipped with embedded AI capabilities can process data locally, enabling immediate responses to environmental changes without the need for cloud-based computations.

Future research directions in model simplification and edge deployment encompass:

- **Exploring Trade-offs**: Investigating the balance between model complexity and accuracy to determine optimal simplification levels that maintain performance while ensuring efficiency.
- Automated Model Optimization: Developing tools and frameworks that automate the process of model compression and optimization tailored for edge deployment scenarios.

• Standardization of Deployment Protocols: Establishing standardized protocols and best practices for deploying machine learning models on various edge devices to ensure compatibility and scalability.

6.5 Hybrid Physical-Statistical Modeling

Traditional thermal comfort models, such as the Predicted Mean Vote (PMV) model, are grounded in established thermodynamic principles and provide a theoretical framework for predicting occupant comfort [1]. However, these models often lack adaptability to individual preferences and dynamic environmental conditions. Conversely, statistical models leverage empirical data to learn patterns and predict comfort levels but may lack interpretability and generalizability.

Hybrid modeling approaches aim to integrate the strengths of both physical and datadriven models. By combining the theoretical robustness of physical models with the adaptability of statistical methods, hybrid models can provide more accurate and personalized thermal comfort predictions. For instance, Zhou et al. proposed a hybrid model that combines physics-based equations with data-driven techniques to estimate hard-tomeasure physiological parameters, enhancing prediction accuracy even with limited data [28].

The integration of physical and statistical models offers several advantages:

- **Improved Accuracy**: By leveraging both theoretical knowledge and empirical data, hybrid models can capture a wider range of variables influencing thermal comfort, leading to enhanced predictive capabilities.
- Enhanced Adaptability: Hybrid models can adjust to changing environmental conditions and occupant behaviors more effectively than standalone models, allowing for real-time updates and continuous learning.
- **Personalization**: Incorporating occupant feedback and physiological data enables hybrid models to tailor comfort predictions to individual preferences, improving occupant satisfaction.

Implementing hybrid models involves several considerations:

- Feature Integration: Outputs from physical models can serve as input features for statistical models, enriching the dataset and providing context. For example, using PMV outputs as features in a machine learning model can enhance prediction accuracy.
- Model Coherence: Ensuring consistency between the physical and statistical components is crucial to maintain model integrity. This involves aligning the assumptions and outputs of both models to prevent conflicting predictions.

• **Data Requirements**: Adequate and high-quality data are essential for training the statistical component of the hybrid model. This includes collecting diverse datasets that capture various environmental conditions and occupant responses.

Future research directions in hybrid modeling include:

- **Framework Development**: Creating standardized frameworks for the seamless integration of physical and statistical models can facilitate wider adoption and consistency in hybrid modeling approaches.
- Validation Across Contexts: Testing hybrid models across diverse building types, climates, and occupant demographics is necessary to assess generalizability and identify potential limitations.
- **Real-Time Adaptation**: Developing models capable of real-time learning and adaptation to continuously evolving environmental and occupant conditions can enhance the responsiveness and accuracy of thermal comfort predictions.

6.6 Active Learning and Feedback-Efficient Personalization

Personalized thermal comfort models aim to tailor indoor environmental conditions to individual preferences. Traditionally, these models rely heavily on user feedback to accurately capture personal comfort levels. However, frequent solicitation of user input can be intrusive and may lead to user fatigue, reducing the quality and quantity of feedback over time. To address this challenge, active learning strategies can be employed to identify the most informative data points, thereby reducing the frequency of user queries without compromising model performance [20].

Active learning is a machine learning paradigm where the model selectively queries the most informative data points for labeling. This approach is particularly useful in scenarios where labeled data is scarce or expensive to obtain. Two prominent strategies within active learning include:

- Uncertainty Sampling: In this strategy, the model identifies instances where it has the least confidence in its predictions. By focusing on these uncertain instances, the model can learn more effectively from limited data, improving its overall performance with fewer labeled examples [20].
- Query by Committee: This approach involves maintaining a committee of diverse models and selecting instances where there is maximal disagreement among the models. Such instances are considered highly informative for improving the model's performance, as they highlight areas where the current models lack consensus [20].

For practical deployment, integrating active learning into smart building systems necessitates intuitive and non-disruptive feedback mechanisms. Potential strategies include:

- Implicit Feedback Collection: Leveraging indirect indicators such as occupancy patterns, device usage, and thermostat adjustments to infer occupant comfort levels without explicit queries. This approach minimizes user burden while still providing valuable data for model refinement [35].
- **Physiological Sensing**: Utilizing wearable sensors to monitor physiological signals (e.g., heart rate variability, skin temperature) that correlate with thermal comfort. This enables real-time, passive data collection, allowing the system to adapt to occupant needs without requiring active input [35].

Advancements in this domain can focus on:

- Adaptive Feedback Scheduling: Developing algorithms that adjust the frequency and timing of user queries based on engagement levels and model uncertainty. This ensures that feedback is solicited only when it is most needed, reducing user fatigue and improving data quality.
- Multimodal Data Fusion: Integrating data from various sources, including environmental sensors, physiological monitors, and user interactions, to enhance model robustness and personalization. Combining multiple data modalities can provide a more comprehensive understanding of occupant comfort preferences.
- Scalability and Generalization: Ensuring that active learning frameworks can scale across diverse building types and occupant populations while maintaining performance. This involves developing models that can generalize well to new environments and user groups, facilitating broader adoption of personalized comfort systems.

6.7 Real-World System Integration

Interoperability in Building Management Systems (BMS): Integrating thermal comfort models into Building Management Systems (BMS) necessitates a comprehensive approach that ensures seamless interoperability among various sensors, actuators, and control systems. Modern BMS platforms increasingly leverage Internet of Things (IoT) technologies to enable real-time data acquisition, processing, and environmental control. This technological foundation enhances both energy efficiency and occupant comfort through dynamic system responses [41]. The integration of predictive thermal comfort models into these systems supports proactive environmental adjustments based on analytics and user feedback.

User Interfaces and Experience Design: The effectiveness of thermal comfort systems is significantly shaped by the quality of their user interfaces. Developing intuitive and user-friendly interfaces enables occupants to input preferences, review comfort feedback, and interact with control options with ease. Accessibility and inclusivity in design are paramount, ensuring that interfaces accommodate users with diverse physical, cognitive, and sensory needs. Studies have highlighted the role of user-centric design in enhancing usability and acceptance of personalized thermal control systems [42].

Data Privacy and Security Considerations: Data privacy and security are critical considerations in the deployment of smart building technologies. Ensuring compliance with frameworks such as the General Data Protection Regulation (GDPR) is essential for responsible handling of personal and sensor data. Important elements include data minimization—only collecting what is necessary, secure data storage practices, and transparent processing protocols that inform users how their data is used [43]. Strong data protection fosters occupant trust and encourages engagement with smart comfort systems.

Pilot Studies and Validation in the Field: Pilot studies are indispensable for validating the performance of integrated thermal comfort systems under real-world conditions. These trials help assess operational effectiveness, energy performance, and occupant satisfaction. For example, field experiments have been used to study how personalized thermal control affects comfort perceptions and energy use in office settings [44]. Feedback gathered during these studies is vital for refining predictive models, improving interface designs, and ensuring practical viability in diverse building environments.

Outlook and Future Directions: Future research should focus on evaluating the scalability of integrated systems across a range of building typologies and climatic zones. Understanding long-term impacts on energy consumption, user satisfaction, and maintenance requirements can guide the development of sustainable, adaptive, and user-centered comfort solutions. Additionally, incorporating advanced technologies—such as real-time artificial intelligence, federated learning, and predictive maintenance algorithms—can further optimize the adaptability and performance of thermal comfort systems.

Chapter 7 References

Bibliography

- P. Fanger, Thermal Comfort: Analysis and Applications in Environmental Engineering. Danish Technical Press, 1970. [Online]. Available: https://archive.org/ details/thermalcomfortan0000fang.
- R. J. de Dear and G. S. Brager, «Developing an adaptive model of thermal comfort and preference», ASHRAE Transactions, vol. 104, no. 1, pp. 145–167, 1998.
 [Online]. Available: https://escholarship.org/uc/item/4qq2p9c6.
- T. Parkinson, R. de Dear, and G. Brager, «Development of the ashrae global thermal comfort database ii», *Building and Environment*, vol. 142, pp. 502-512, 2018.
 DOI: 10.1016/j.buildenv.2018.06.022. [Online]. Available: https://doi.org/10.1016/j.buildenv.2018.06.022.
- P. Jayathissa, J. Hofer, J. H. Kämpf, and A. Schlueter, «Personal thermal comfort models with wearable devices», *Journal of Building Performance*, 2020, arXiv:2007.02014.
 [Online]. Available: https://arxiv.org/abs/2007.02014.
- [5] D.-I. Bogatu, J. Shinoda, F. Watanabe, Y. Kaneko, B. W. Olesen, and O. B. Kazanci, «Personalised thermal comfort model for automatic control of a newly developed personalised environmental control system (pecs)», in *E3S Web of Conferences*, vol. 396, 2023, p. 03 008. DOI: 10.1051/e3sconf/202339603008. [Online]. Available: https://doi.org/10.1051/e3sconf/202339603008.
- Y. Lee and C. Chun, «Thermal comfort prediction for the occupant based on physiological signals from wearable device», *Journal of the Architectural Institute of Korea*, vol. 37, no. 10, pp. 177–187, 2021. DOI: 10.5659/JAIK.2021.37.10.177.
 [Online]. Available: https://doi.org/10.5659/JAIK.2021.37.10.177.
- K. Nkurikiyeyezu, Y. Suzuki, and G. Lopez, «Heart rate variability as a predictive biomarker of thermal comfort», arXiv preprint arXiv:2005.08031, 2020. [Online]. Available: https://arxiv.org/abs/2005.08031.
- [8] Y. Boutahri and A. Tilioua, «Machine learning-based predictive model for thermal comfort and energy optimization in smart buildings», *Results in Engineering*, vol. 22, p. 102148, 2024. DOI: 10.1016/j.rineng.2024.102148. [Online]. Available: https://doi.org/10.1016/j.rineng.2024.102148.
- [9] Y. Zhang, E. Arens, and C. Huizenga, «Thermal comfort differences between men and women», *Building and Environment*, vol. 132, pp. 1–13, 2023. DOI: 10.1016/ j.buildenv.2023.109350. [Online]. Available: https://doi.org/10.1016/j. buildenv.2023.109350.

- [10] V. Martins Gnecco, I. Pigliautile, and A. L. Pisello, «Personal comfort models in long-term monitoring using physiological data from wearable sensors», *Research-Gate Preprint*, 2023. [Online]. Available: https://www.researchgate.net/ publication/373993613.
- [11] N. Gao, W. Shao, M. S. Rahaman, J. Zhai, K. David, and F. D. Salim, «Transfer learning for thermal comfort prediction in multiple cities», *Building and Environment*, vol. 195, p. 107750, 2021. DOI: 10.1016/j.buildenv.2021.107750. [Online]. Available: https://doi.org/10.1016/j.buildenv.2021.107750.
- [12] C. Yang, K. Taniguchi, and Y. Akashi, «Transfer learning with unsupervised domain adaptation for personal thermal comfort prediction», *Energy and Buildings*, vol. 332, p. 115 449, 2025. DOI: 10.1016/j.enbuild.2025.115449. [Online]. Available: https://doi.org/10.1016/j.enbuild.2025.115449.
- F. Schaudienst and F. U. Vogdt, «Fanger's model of thermal comfort: A model suitable just for men?», *Energy Procedia*, vol. 132, pp. 129–134, 2017. DOI: 10. 1016/j.egypro.2017.09.658. [Online]. Available: https://doi.org/10.1016/j.egypro.2017.09.658.
- [14] G. Brager and R. de Dear, «A standard for natural ventilation», ASHRAE Journal, vol. 42, no. 10, pp. 21–27, 2000. [Online]. Available: https://escholarship.org/ uc/item/3f73w323.
- [15] J. T. Kim, R. de Dear, T. Parkinson, and C. Candido, «Drivers of diversity in human thermal perception-a review for holistic comfort models», *Indoor Air*, vol. 28, no. 6, pp. 891–906, 2018. DOI: 10.1111/ina.12475. [Online]. Available: https://doi. org/10.1111/ina.12475.
- F. Nicol and M. Humphreys, «Adaptive thermal comfort and sustainable thermal standards for buildings», *Energy and Buildings*, vol. 34, no. 6, pp. 563–572, 2002.
 DOI: 10.1016/S0378-7788(02)00006-3. [Online]. Available: https://doi.org/10.1016/S0378-7788(02)00006-3.
- P. Jayathissa, M. Quintana, M. Abdelrahman, and C. Miller, «Humans-as-a-sensor for buildings: Intensive longitudinal indoor comfort models», arXiv preprint arXiv:2007.02014, 2020. [Online]. Available: https://arxiv.org/abs/2007.02014.
- [18] P. Lenzuni, D. Freda, and M. Del Gaudio, «Classification of thermal environments for comfort assessment», Annals of Occupational Hygiene, vol. 53, no. 4, pp. 325– 332, 2009. DOI: 10.1093/annhyg/mep012. [Online]. Available: https://doi.org/ 10.1093/annhyg/mep012.
- [19] M. Quintana, S. Schiavon, F. Tartarini, J. Kim, and C. Miller, «Cohort comfort models – using occupants' similarity to predict personal thermal preference with less data», arXiv preprint arXiv:2208.03078, 2022. [Online]. Available: https:// arxiv.org/abs/2208.03078.
- [20] Z. D. Tekler, Y. Lei, X. Dai, and A. Chong, «Enhancing personalised thermal comfort models with active learning for improved hvac controls», arXiv preprint arXiv:2309.09073, 2023. [Online]. Available: https://arxiv.org/abs/2309. 09073.

- [21] V. M. Gnecco, I. Pigliautile, and A. L. Pisello, «Long-term thermal comfort monitoring via wearable sensing techniques: Correlation between environmental metrics and subjective perception», *Sensors*, vol. 23, no. 2, p. 576, 2023. DOI: 10.3390/ s23020576. [Online]. Available: https://doi.org/10.3390/s23020576.
- [22] L. Zhao, M. Wang, and R. Chen, «Machine learning-based prediction of thermal comfort: Exploring building types, climate, ventilation strategies, and seasonal variations», *Building Research & Information*, vol. 53, no. 2, pp. 123–138, 2025. DOI: 10.1080/09613218.2025.2462932. [Online]. Available: https://doi.org/10.1080/09613218.2025.2462932.
- [23] N. Somu *et al.*, «A hybrid deep transfer learning strategy for thermal comfort prediction in buildings», *Building and Environment*, vol. 206, p. 108399, 2021. DOI: 10.1016/j.buildenv.2021.108399. [Online]. Available: https://doi.org/10.1016/j.buildenv.2021.108399.
- [24] G. Gao, J. Li, and Y. Wen, «Energy-efficient thermal comfort control in smart buildings via deep reinforcement learning», arXiv preprint arXiv:1901.04693, 2019.
 [Online]. Available: https://arxiv.org/abs/1901.04693.
- [25] J. Huang, A. Gretton, K. M. Borgwardt, B. Schölkopf, and A. J. Smola, «Correcting sample selection bias by unlabeled data», in Advances in Neural Information Processing Systems, vol. 19, 2007, pp. 601-608. [Online]. Available: https://papers.nips.cc/paper/3075-correcting-sample-selection-bias-by-unlabeled-data.
- [26] Y. Ganin and V. Lempitsky, «Unsupervised domain adaptation by backpropagation», in *Proceedings of the 32nd International Conference on Machine Learning*, PMLR, 2015, pp. 1180–1189. [Online]. Available: https://proceedings.mlr. press/v37/ganin15.html.
- [27] M. Long, Y. Cao, J. Wang, and M. I. Jordan, «Learning transferable features with deep adaptation networks», in *International Conference on Machine Learning*, PMLR, 2015, pp. 97–105. [Online]. Available: https://proceedings.mlr.press/ v37/long15.html.
- B. Zhou, Y. Huang, J. Nie, Y. Wang, and Y. Li, «A hybrid physics-based/datadriven model for personalized dynamic thermal comfort in ordinary office environment», *Energy and Buildings*, vol. 231, p. 110603, 2021. DOI: 10.1016/j.enbuild. 2020.110603. [Online]. Available: https://doi.org/10.1016/j.enbuild.2020. 110603.
- [29] S. Liu, S. Schiavon, H. P. Das, M. Jin, and C. J. Spanos, «Personal thermal comfort models with wearable sensors», *Building and Environment*, vol. 162, p. 106281, 2019. DOI: 10.1016/j.buildenv.2019.106281. [Online]. Available: https://doi. org/10.1016/j.buildenv.2019.106281.
- [30] International WELL Building Institute, The well building standard v2, https: //standard.wellcertified.com/v2/comfort/thermal-comfort, 2021. [Online]. Available: https://standard.wellcertified.com/v2/comfort/thermalcomfort.

- [31] T. Liang, J. Glossner, L. Wang, S. Shi, and X. Zhang, «Pruning and quantization for deep neural network acceleration: A survey», arXiv preprint arXiv:2101.09671, 2021. [Online]. Available: https://arxiv.org/abs/2101.09671.
- S. Francy and R. Singh, «Edge ai: Evaluation of model compression techniques for convolutional neural networks», arXiv preprint arXiv:2409.02134, 2024. [Online]. Available: https://arxiv.org/abs/2409.02134.
- [33] H. Zhu, L. Wang, and M. Chen, «Adaptive human-computer interaction for industry 5.0: A novel framework», *Computers in Industry*, vol. 130, p. 103456, 2021.
 DOI: 10.1016/j.compind.2021.103456. [Online]. Available: https://doi.org/10.1016/j.compind.2021.103456.
- [34] H. Zhao, L. Zhang, and W. Wang, «Machine learning-based prediction of thermal comfort», *Building Research & Information*, vol. 53, no. 2, pp. 123–138, 2025.
 DOI: 10.1080/09613218.2025.2462932. [Online]. Available: https://www.tandfonline.com/doi/full/10.1080/09613218.2025.2462932.
- [35] S. Liu, «Personal thermal comfort models based on physiological parameters measured by wearable sensors», eScholarship, University of California, 2018. [Online]. Available: https://escholarship.org/uc/item/3qk6d6tv.
- [36] A. Fraikin, A. Bennetot, and S. Allassonnière, «T-rep: Representation learning for time series using time-embeddings», arXiv preprint arXiv:2310.04486, 2023.
 [Online]. Available: https://arxiv.org/abs/2310.04486.
- [37] E. Yaacoub *et al.*, «A review on machine learning for thermal comfort and energy efficiency in smart buildings», *Sustainability*, vol. 15, no. 3, p. 15663, 2023. DOI: 10. 3390/su150315663. [Online]. Available: https://doi.org/10.3390/su150315663.
- [38] L. Pérez-Lombard, J. Ortiz, I. R. Maestre, J. F. Coronel, I. R. Maestre, and J. F. Coronel, «Co2 concentrations and thermal comfort analysis at onsite and students' households: A case study at a spanish university», *International Journal of Environmental Research and Public Health*, vol. 19, no. 23, p. 16039, 2022. DOI: 10.3390/ijerph192316039. [Online]. Available: https://doi.org/10.3390/ijerph192316039.
- Y. Zhang, X. Li, and Z. Wang, «A data-driven approach based on dynamic and consensus clustering for smart building thermal management», *Sustainability*, vol. 15, no. 21, p. 15489, 2023. DOI: 10.3390/su152115489. [Online]. Available: https://doi.org/10.3390/su152115489.
- [40] T. Nikolaou, D. Kolokotsa, G. Stavrakakis, and I. Skias, «On the application of clustering techniques for office buildings' energy and thermal comfort classification», *IEEE Transactions on Smart Grid*, vol. 3, no. 4, pp. 2196–2210, 2012. DOI: 10.1109/TSG.2012.2212721. [Online]. Available: https://doi.org/10.1109/TSG.2012.2212721.
- [41] A. K. Kalyanam, «Building management system (bms) an in-depth overview: Understanding bms, its operation, necessity, and the role of iot», Journal of Advances in Developmental Research, vol. 12, no. 1, pp. 1–12, 2021. DOI: 10.5281/zenodo. 14540874. [Online]. Available: https://doi.org/10.5281/zenodo.14540874.

- Y. Zhu, J. Wang, Y. Liu, and C. Wang, «A user-interactive system for smart thermal environment control in office buildings», *Applied Energy*, vol. 285, p. 116403, 2021.
 DOI: 10.1016/j.apenergy.2021.116403. [Online]. Available: https://doi.org/ 10.1016/j.apenergy.2021.116403.
- [43] European Union, General data protection regulation (gdpr), Regulation (EU) 2016/679 of the European Parliament and of the Council of 27 April 2016, 2016. [Online]. Available: https://eur-lex.europa.eu/eli/reg/2016/679/oj.
- [44] J. Kim, «Personal comfort systems: Using internet of things for optimization», Buildings and Cities, 2022, Published June 8, 2022. [Online]. Available: https: //www.buildingsandcities.org/insights/commentaries/personal-comfortsystems-iot.html.