Energy & Buildings 319 (2024) 114507

Contents lists available at ScienceDirect

Energy & Buildings

journal homepage: www.elsevier.com/locate/enbuild

Data-efficient comfort modeling: Active transfer learning for predicting personal thermal comfort using limited data

Zeynep Duygu Tekler, Yue Lei, Adrian Chong *

Department of the Built Environment, National University of Singapore, 4 Architecture Drive, 117566, Singapore

ARTICLE INFO

Dataset link: https://github.com/ideas-lab-nus/ComfortPredict-ActiveTL

Keywords:
Data-efficient machine learning
Active learning
Transfer learning
Personal thermal comfort models
CNNLSTM

ABSTRACT

Personal comfort models are used to predict thermal comfort responses at the individual level rather than predicting the average thermal comfort responses for large populations. These models, data-driven in nature, need to be trained on large amounts of occupant comfort feedback and sensor data to achieve accurate predictions. However, collecting such data is often expensive and labor-intensive in reality. To address this, we proposed a data-efficient active transfer learning (ATL) framework to improve the performance of personal comfort models under limited data. To demonstrate the validity of this framework, we developed a base Convolutional Neural Network-Long Short-term Memory (CNNLSTM) model alongside two transfer learning models utilizing feature extraction (TL-CNNLSTM-FE) and fine-tuning (TL-CNNLSTM-FE) approaches, enhanced by a novel active learning strategy. Using these models, three comfort prediction tasks (i.e., thermal preference, thermal acceptability, and air movement preference) were performed by transferring the knowledge from the ASHRAE Global Thermal Comfort Database II to a limited dataset collected in the tropics. Empirical results indicate that the active transfer learning framework proposed was able to consistently outperform the base and transfer learning models using only less than 10% of the training data for all personal comfort tasks, highlighting the effectiveness of this strategy. The implications of this work are especially useful for the research community working on the practical applications of data-efficient machine learning approaches for personal thermal comfort predictions.

1. Introduction

There is a growing trend in learning and predicting individual thermal comfort needs through data gathered from everyday environments. Personalized thermal comfort models play a pivotal role in this shift, aiming to predict thermal comfort responses at the individual level rather than relying on the average response of large populations. Unlike traditional models, including Predictive Mean Vote (PMV) and adaptive models [1], personal comfort models are designed to account for unique differences between each occupant, allowing for a more profound comprehension of their distinct thermal comfort needs in the buildings. Such information not only enhances occupants' comfort but also contributes to energy efficiency by ensuring that the building systems operate optimally to meet specific comfort requirements [2]. A notable example is occupant-centric control applications, which integrate occupants' personal comfort and preferences in building controls [3]. These include personalized lighting control [4] to improve visual satisfaction and lighting energy use, personalized plug load control [5] based on occupancy and personalized plug load use preferences, as well as occupant-centric Heating, Ventilation and Air Conditioning (HVAC) controls [6] based personal thermal comfort and occupancy to improve the efficiency and effectiveness of HVAC systems. Past studies suggest that smart HVAC systems can improve energy savings and thermal comfort by more than 15% [7], while employing an appropriate comfort model for HVAC controls can lead to further energy savings exceeding 10% during simulations [8]. Meanwhile, in real-world implementations, personal preference-based HVAC control systems have demonstrated remarkable energy savings, ranging from 28% to 35% [9], proving the effectiveness of personal comfort models in improving energy efficiency.

Personal comfort models, being inherently data-driven machine learning models, require two primary categories of data for their development. (1) The first category includes objective measurements, including indoor and outdoor environmental conditions, HVAC operations, building occupancy, and occupants' physiological conditions. These measurements are obtained through various sensing and communication technologies. (2) The second category involves collecting occupant comfort feedback through subjective surveys. This feedback

* Corresponding author.
E-mail addresses: zdtekler@nus.edu.sg (Z.D. Tekler), bdgleiy@nus.edu.sg (Y. Lei), adrian.chong@nus.edu.sg (A. Chong).

https://doi.org/10.1016/j.enbuild.2024.114507
Received 27 February 2024; Received in revised form 25 May 2024; Accepted 29 June 2024
Available online 3 July 2024

indicates occupants' thermal preference, thermal acceptability, or thermal sensation based on the thermal conditions they experience.

While recent developments in the Internet of Things (IoT) and sensing technologies have enabled the cost-efficient and effective collection of various sensor measurements, challenges persist during the collection of occupant comfort feedback, also known as user-labeled data. These challenges are primarily due to the intrusive and labor-intensive nature of the data collection process, often leading to survey fatigue during the data labeling. Researchers also reported difficulties determining an appropriate frequency to collect comfort feedback from the study participants [10]. A recent study further quantified this challenge, finding that approximately 250–300 data points per study participant are necessary for achieving accurate predictions of personal thermal comfort, highlighting the significant volume of data needed [11]. Therefore, it is essential to strike a balance between gathering sufficient data to develop accurate personal comfort models while minimizing the effort required from respondents. This challenge highlights the crucial need for developing more data-efficient approaches that can accurately predict personal comfort with less reliance on extensive comfort feedback.

In this paper, our objective is to reduce the need of collecting large sums of sensor and occupant comfort survey data to develop data-efficient personal comfort models with satisfactory model performance.

The study contributions are outlined as follows:

- Introduced a novel active transfer learning framework that integrates active learning and transfer learning approaches to improve the performance of personal comfort predictions using limited data.
- Presented a comprehensive comfort evaluation for three comfort tasks (i.e., thermal acceptability, thermal preference, and air movement preference) by transferring the knowledge from one of the largest global thermal comfort databases (i.e., ASHRAE Global Thermal Comfort Database II) to our limited dataset collected in the tropics.
- Developed two transfer learning approaches (i.e., feature extraction and fine-tuning) based on a CNNLSTM architecture and evaluated them with different amounts of training data under air-conditioned (AC) and natural ventilation (NV) conditions.

The remainder of this paper is structured as follows: Section 2 reviews the existing personal comfort models and data-efficient machine learning approaches for personal comfort models. Section 3 provides details on the active transfer learning methodology, including the utilized datasets, the architecture of the base CNNLSTM model, transfer learning approaches, and the novel active learning strategy developed. Section 4 presents the results and analysis, evaluating the impact of active transfer learning strategy and different amounts of training data on the performance of personal comfort models. Finally, Section 5 concludes the paper by summarizing the key findings of this research.

2. Related work

2.1. Development of personal comfort models

In recent years, there has been a notable rise in the popularity of data-driven thermal comfort modeling, leading to many efforts leveraging machine learning (ML) techniques for predicting personal thermal comfort. Popular machine learning models, such as gradient boosting [12], random forest (RF) classification [13], logistic regression (LR) [14], support vector machine (SVM) [15], and have been employed to predict personal thermal comfort based on different thermal metrics. Notably, Luo et al. [12] performed a comprehensive comparison of nine common machine learning algorithms for predicting thermal sensation using the ASHRAE Comfort Database II. Their findings indicated that machine learning models exhibited superior accuracy compared to traditional PMV models, with Random Forest outperforming other algorithms. This observation particularly aligns with Kim et al.'s review of different models [10], emphasizing a significant improvement in predictive accuracy (17-40%) when compared with conventional comfort models (i.e., PMV and adaptive models). This underscores the need for adopting more personalized approaches for predicting thermal comfort.

The data used for developing personal comfort models were obtained through (1) various sensor measurements and (2) occupant comfort feedback via surveys. Sensor measurements can be categorized under five broad categories: indoor environmental measurements, outdoor weather measurements, occupant-related measurements, temporal measurements, and system (i.e., HVAC) related measurements. Indoor environmental measurements include indoor temperature (T_{ia}), indoor relative humidity (RH_i), indoor (CO_2) levels, indoor radiant temperature (T_r), indoor air velocity (AV_i). Outdoor weather measurements often involve collecting outdoor humidity, outdoor air temperature (T_{oa}), solar radiation, wind speed, and wind direction. Both indoor and environmental measurements have been widely used to predict thermal preferences [16], thermal sensations [17], and thermal acceptability [18]. Occupant-related measurements encompass both physiological and behavioral data. To incorporate occupant physiological measurements into thermal comfort modeling, specialized sensors are utilized to capture parameters such as skin temperatures, heart rate, and metabolic rate [19]. Furthermore, occupant activity levels (i.e., exercise habits) and clothing levels are collected as part of behavioral measurements [20]. Temporal measurements, such as the day of the week or specific hours of the day, represent the variations over time in the thermal comfort levels experienced by occupants Finally, HVAC-related measurements represent the operational status of indoor environments, encompassing variables such as damper position, supply airflow rate, and variable air volume control settings [10].

In addition to sensor measurements, the development of personal thermal comfort models necessitates gathering feedback on occupants' comfort using various methods. Surveys are frequently employed, utilizing a range of interfaces such as wearables or thermostats to collect comfort feedback efficiently [21]. Occupants' comfort feedback can be collected based on their thermal acceptability (Acceptable, Unacceptable), thermal preference (Cooler, No change, Warmer), and thermal sensation (7-point scale Predicted Mean Vote (PMV) index). Furthermore, air movement preferences (More, No Change, Less) and air movement acceptability (Acceptable, Unacceptable) were also considered as comfort measures for mixed-mode buildings in recent studies [22].

2.2. Active learning and transfer learning approaches for personal comfort models

Data-efficient machine learning approaches have been increasingly gaining attention for personal comfort modeling, addressing the challenges associated with collecting large amounts of training data to achieve accurate personal comfort models [23]. Among the small number of studies available, these approaches enhance the performance of traditional machine-learning models when dealing with limited training data. This is especially useful as collecting frequent occupant comfort feedback (i.e., user-labeled data) is expensive, labor-intensive, and often impractical for extended data collection periods.

Active learning (AL) is a promising approach for predicting personal comfort with limited training data. AL is used to identify the most informative data points for user labeling, thereby potentially reducing the frequency of user comfort feedback collection. A pioneering study [24] demonstrated that AL could reduce user labeling efforts by up to 46% in personalized thermal comfort models during a field implementation. The same study also implemented a feature selection algorithm [25] to understand and filter the useful input features before applying AL as part of a data-efficient minimum sensing strategy. The application of AL has been extended to HVAC control systems, with results indicating a substantial decrease in labeling effort (about 31.0%) while still maintaining high levels of personal comfort and energy savings compared to conventional controls [26].

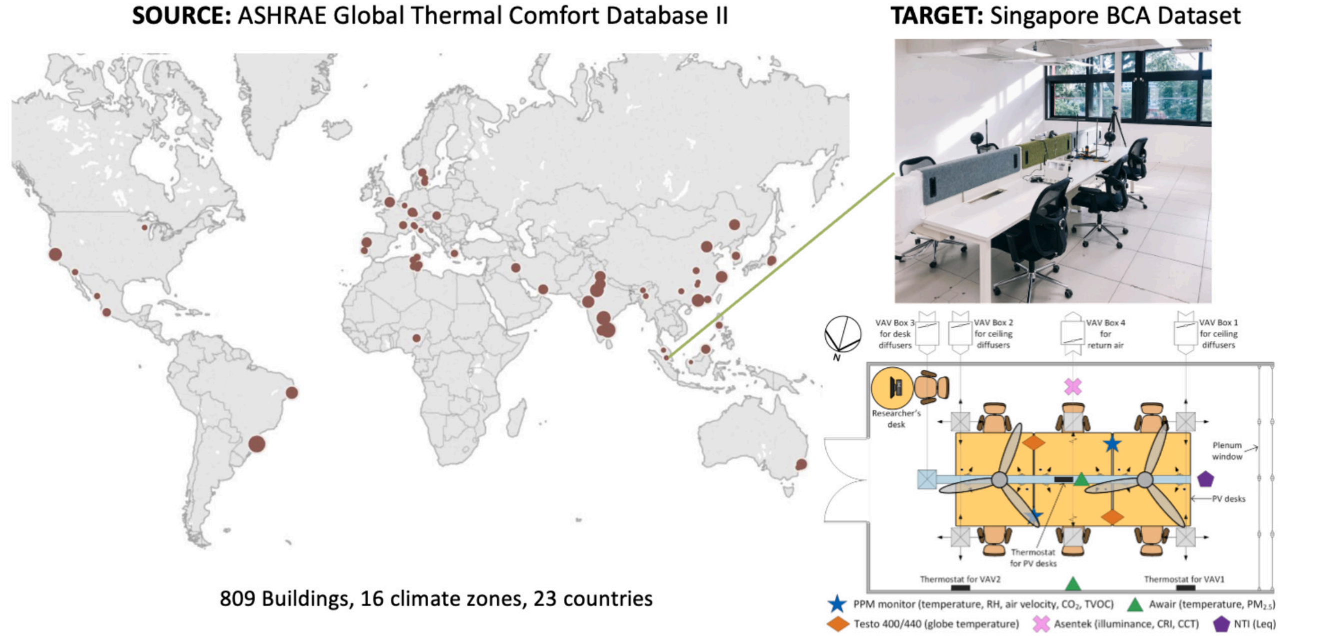

Fig. 1. The ASHRAE Global Thermal Comfort Database II (left) and Singapore BCA Dataset (right) serve as the source and target domains in this study, respectively.

Transfer Learning (TL) presents an alternative strategy involving the application of pre-trained models from an information-rich (source) domain to a related task in a different (target) domain where labeled data is limited or unavailable. Gao et al. [27] investigated TL for predicting thermal comfort across different cities. They utilized multilayer perception (MLP) models trained on data from various climatic regions and applied them to a (target) building containing scarce data. Their findings indicated that models pre-trained on comparable climate zones demonstrated improved prediction accuracy. Similarly, Somu et al. [28] developed a transfer learning-based Convolutional Neural Networks-Long Short Term Memory (CNNLSTM) model for thermal preference prediction. The resulting model reported 56% accuracy, using the same source and target domains with [27]. Furthermore, Park et al. [29] introduced an ensemble TL approach, leveraging knowledge from diverse thermal conditions and physiological data to predict thermal comfort preferences accurately for specific subjects with sparse data. Finally, Li et al. [30] proposed a transfer learning model using TrAdaBoost and introduced an automatic weighting strategy to adjust the weights of the training data from the source domain for performance improvement. To date, only one study [31] has attempted to combine both active and transfer learning for predicting personal thermal comfort responses. The study involved using a ridge regression model to learn the thermal comfort distribution of a group of occupants within a building before transferring the model to a different group of occupants from the same building by leveraging limited queries. However, the study is limited as it still requires collecting large amounts of comfort feedback data from other occupants in the building initially before it can be transferred to a new occupant within the same building, significantly reducing the approach's scalability when applied at a large scale.

Our study extends upon past studies by proposing a novel active transfer learning framework that leverages one of the largest global comfort datasets, the ASHRAE Global Thermal Comfort Database II, to train the initial model on the personal comfort preferences of occupants from different buildings. The model is subsequently transferred to a new group of occupants using a limited comfort dataset collected from the new building. This proposed method significantly improves the approach's scalability when applied at a large scale. Additionally, we further demonstrated the robustness and effectiveness of the approach by being the first study to perform a comprehensive evaluation of different active transfer learning strategies based on a diverse range of comfort tasks (i.e., thermal preference, thermal acceptability, and air movement preference) using a real-world comfort dataset collected in the tropics.

3. Methodology

3.1. Dataset description

Two thermal comfort datasets were utilized for this study. The first dataset, the ASHRAE Global Thermal Comfort Database II [32], serves as the information-rich source domain, containing an extensive amount of thermal comfort data collected from multiple buildings across different climatic zones. The second dataset, the Singapore Building Construction Authority (BCA) testbed dataset [22], represents the information-poor target domain, consisting of a limited amount of thermal comfort data collected from a single building in a tropical climate. Further elaboration on the data collection process and the fields present in both datasets is provided in the subsequent subsections and depicted in Fig. 1.

3.1.1. Source domain: ASHRAE global thermal comfort database II

The ASHRAE Global Thermal Comfort Database II currently stands as one of the most comprehensive open-source repositories for thermal comfort research, encompassing a compilation of 34 field-measured thermal comfort studies [32]. This database comprises data collected from over 800 commercial and residential buildings, with contributions from researchers representing 23 countries and 39 institutions across 16 different climate zones. Each dataset incorporated into the database undergoes a rigorous quality assurance process to ensure its suitability for robust hypothesis testing. Various steps are taken to maintain quality control, including visualizing field variables to identify outliers, generating cross-plots between highly correlated variables (e.g., thermal comfort and thermal sensation) to detect inaccurately coded data, and randomly sampling from each study to ensure consistency between the original and standardized database.

In cases where certain data fields were absent from the dataset, those fields were replaced with null values to prevent misrepresentation. Through this process, a total of 81,846 data entries, containing information about the building's indoor and outdoor environmental conditions, as well as study participants' demographic information, occupancy patterns, and subjective comfort votes (i.e., thermal acceptability, thermal

Table 1
Summary of operating conditions in the Singapore BCA Dataset.

Mode	Description	Operating Condition	Room Temperature Setpoint (°C)	Ceiling Fan Speed (m/s)
AC	Activating the ambient VAV system with traditional ceiling diffusers	1	24	Off (0.15)
		2	24	Off (0.15)
		3	26	Low (0.45)
		4	26	Medium (0.90)
NV	HVAC System switched off with the plenum windows fully open	5	Free-Floating	Off (0.15)
		6		Low (0.45)
		7		Medium (0.90)
		8		High (1.15)

comfort, thermal preference, air movement preference, and air movement acceptability), were retained. Additionally, the data from the original ASHRAE RP-884 database is also included, resulting in 107,463 total entries.

3.1.2. Target domain: Singapore BCA testbed dataset

The BCA testbed is a 50 m^2 experimental facility consisting of six desks, an extra desk for a researcher, and an operable window located on the facade's west side. The operable window allows the testbed to operate as an air-conditioned or naturally ventilated space. The testbed is conditioned by a variable air volume (VAV) system and a stand-alone air-handling unit with its own building management system for greater precision in controlling space conditions. Additionally, the indoor environmental conditions are monitored by a network of advanced sensors collecting information about indoor relative humidity, indoor air temperature, air velocity, globe temperature, carbon dioxide levels, total volatile organic compounds (TVOC), and indoor fine particulate matter ($PM_{2.5}$). Real-time measurements of outdoor weather conditions are obtained through a weather station, including outdoor relative humidity, outdoor air temperature, outdoor fine particulate matter (PM_1, $PM_{2.5}$, PM_{10}), and atmospheric pressure.

To collect occupants' subjective comfort votes under different ventilation conditions, a data collection effort was conducted in June 2022 [22]. Fifty-eight tropically acclimatized (residing in Singapore for at least the past three years) participants (50% males and 50% females), aged between 21 and 60, took part in the data collection. Throughout the data collection period, the indoor conditions in the testbed were altered every thirty minutes, following a randomized sequence, to one of eight distinct settings. Half of these settings were categorized under Natural Ventilation (NV) mode, while the remaining fell under Air-Conditioned (AC) mode. At the beginning of each setting, adjustments were made manually to the room temperature setpoint and ceiling fan speed, following the specifications outlined in Table 1, to establish distinct operating conditions. A brief thermal comfort survey was distributed via wearable devices assigned to participants at the onset of each setting, at the 5th and 15th minute intervals, to gather information regarding their thermal acceptability, preference, and air movement preference under the current conditions. Additionally, a comprehensive Indoor Environmental Quality (IEQ) survey was conducted at the 25th-minute mark using the Qualtrics online platform to assess participants' current levels of thermal, air quality, acoustic, and visual comfort. This data collection effort resulted in 1,745 entries collected over the course of the study period.

3.2. Data processing and analysis

In this section, we outline the data processing and analysis steps undertaken in our study.

We initiate the data processing step by retaining features common to both the ASHRAE Global Thermal Comfort Database II and the BCA Testbed Dataset, and those features that offer valuable insights into occupants' thermal and air movement comfort conditions. These features provide information about the building's indoor and outdoor environmental conditions, such as indoor room temperature (T_{ia}), indoor relative humidity (RH_i), air velocity (AV), globe temperature (T_g), outdoor temperature (T_{oa}), outdoor relative humidity (RH_o), and the room's ventilation mode (i.e., air-conditioned or naturally ventilated).

Following this step, we addressed the missing data issue in some of the fields within the ASHRAE Global Thermal Comfort Database II. An imputation step is employed to populate the missing data. Despite various imputation algorithms proposed previously [33], we opted for the MissForest imputation algorithm due to its widespread acceptance and effectiveness in handling missing tabular data [34]. The imputation process begins by focusing on the column with the lowest amount of missing values (referred to as the candidate column) and initially imputing the missing values in other columns using their mean values. Subsequently, an RF model is trained, where the candidate column's missing values are predicted based on the values in the other columns. This process is repeated for all of the columns with missing values, iteratively improving the imputed dataset until the difference between successive imputed datasets becomes negligible.

After imputing the missing data within the ASHRAE Global Thermal Comfort Database II, an analysis was performed on the subjective votes reported by the study participants in both the ASHRAE Global Database II and BCA Testbed Dataset (refer to Fig. 2). By examining the distribution of reported labels in Fig. 2, it can be observed that the study participants in the ASHRAE Global Database II tend to more frequently report that they were satisfied with the current environmental conditions by indicating "Acceptable" (72%) when asked about their thermal acceptability, followed by "Unacceptable" (28%). Similarly, for thermal comfort preferences, 50% of the responses indicated "No Change" followed by "Cooler" at 32.9% and finally "Warmer" at 17.1%. A similar trend is observed when asked about their air movement preferences, with the majority of responses indicating "No Change" (55%), followed by "More" at 36% and "Less" at 9%. This trend is also reflected in the subjective votes reported by the study participants in the BCA testbed, with the majority of responses reporting "No Change" when asked about their air movement and thermal preferences, and "Acceptable" when asked about their thermal acceptability. This imbalance between the comfort votes reported by the study participants may lead to the development of a comfort prediction model that is heavily biased toward the majority class (i.e., No Change or Acceptable), and may fail to correctly identify conditions which are uncomfortable for the occupants. Therefore, we addressed the imbalance in the subjective labels by applying the Synthetic Minority Over-sampling Technique (SMOTE) on the ASHRAE Global Thermal Comfort Database II to rebalance the instances belonging to the minority classes.

3.2.1. Synthetic minority over-sampling technique (SMOTE)

The SMOTE algorithm, first proposed by [35], is an oversampling method designed to address the issues associated with class imbalance by creating synthetic instances from the minority class. The algorithm begins by sampling the k nearest neighbors of a instance in the minority class X_j and creating a new instance P_k through a random interpolation process between X_j and one of the nearest neighbors, denoted as

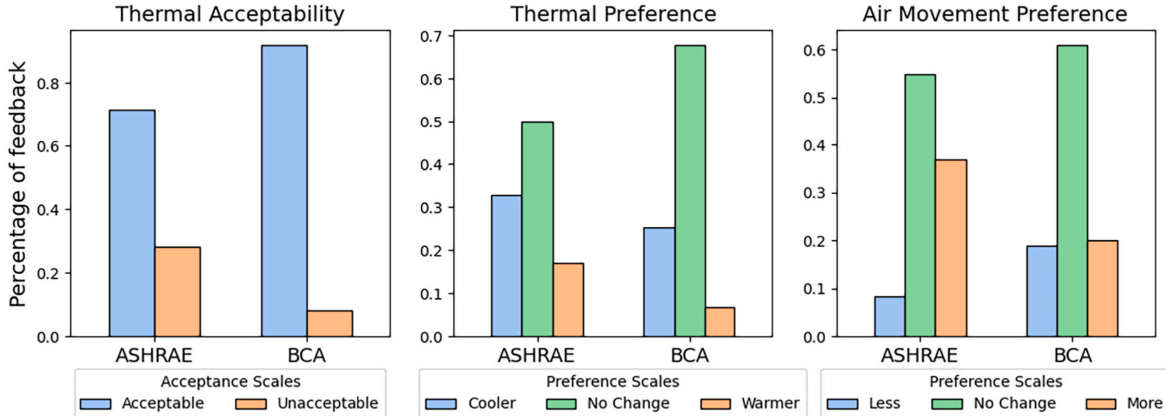

Fig. 2. Label distribution of the comfort feedback for thermal acceptability, thermal preference, and air movement preference reported in the ASHRAE Global Thermal Comfort Database II and the Singapore BCA Dataset. (For interpretation of the colors in the figure(s), the reader is referred to the web version of this article.)

$Z_j k$. This process can be mathematically represented using the following equation [36]:

$$P_k = X_j + random(0,1) * (Z_{kj} - X_j). \tag{1}$$

In essence, the SMOTE algorithm functions by creating new instances along the line segments connecting x_j with the selected neighbor $Z_j k$. Hence, the fundamental assumption underlying the algorithm is that the instances lying between two neighboring instances in the minority class also belong to the same minority class.

Given that the ASHRAE Global Thermal Comfort Database II contains a combination of both nominal and continuous variables, this study adopted the SMOTE-Nominal Continuous algorithm, which extends upon the SMOTE algorithm, to perform oversampling of the minority class within the dataset. The oversampled version of the ASHRAE Global Thermal Comfort Database II is used for model training in the subsequent sections of this study.

3.3. Problem formulation

The task of predicting occupants' thermal preference, thermal acceptability, and air movement preference can be framed as a classification problem. This entails defining a probabilistic function that assigns a probability score $P(y|X)$, indicating the likelihood that a sample belongs to a specific class label y given the sample's input features X. The input feature X, in this case, represents a tuple containing the environmental features, while the class label y depends on whether the model is predicting the occupants' thermal preference $y \in$ No Change, Warmer, Cooler, thermal acceptability $y \in$ Acceptable, Unacceptable, or air movement preference $y \in$ No Change, More, Less.

The probabilistic function $f(.)$ is obtained by learning from a training dataset D_{train} before validating its performance on a test dataset D_{test}. In the case of a regular classification problem, both D_{train} and D_{test} originate from the same domain D. However, in the context of transfer learning, the problem extends to involve two separate domains: the source domain D_{source} and the target domain D_{target}.

The source domain D_{source} is typically defined as the information-rich domain where a large number of labeled instances are readily available, while the information-poor target domain D_{target} tends to contain only a limited number of labeled instances due to the high cost of data collection and human labeling (i.e., $|D_{source}| >> |D_{target}|$). Therefore, the motivation behind transductive transfer learning involves enhancing the predictive performance of the probabilistic function $f(.)$ in the target domain D_{target}, by leveraging on the vast amount of labeled instances in the source domain D_{source}, and applying that knowledge in the target domain D_{target}.

3.4. Base model (CNNLSTM) architecture

The base model architecture chosen for this study is based on the Convolutional Neural Network-Long Short-term Memory (CNNLSTM) model. This decision is due to the model's ability to effectively capture the spatio-temporal relationship between different time-series variables of the building's indoor and outdoor environmental conditions to perform accurate comfort modeling [28]. Additionally, the model also allows us to freeze different parts of the CNN and LSTM layers to facilitate the implementation and evaluation of different transfer learning strategies.

Comprising a convolutional layer, two LSTM layers, two dense layers, and an output layer, the CNNLSTM model architecture is designed to capture the spatiotemporal characteristics of the building's environmental features X to predict the most likely response in terms of the occupants' thermal preference, thermal acceptability, or air movement preference y.

A forward pass into the CNNLSTM model involves passing an input feature tuple X of dimension (7, 1) into a 1D-convolutional layer with a filter size of 128 and kernel size of 5, to result in an output tensor of dimension (7, 128). In this case, the convolutional layer functions as a feature extractor, discerning the spatially invariant structures within input X. To prevent overfitting, a dropout layer with a probability of 0.1 is introduced after the 1D-convolutional layer. The spatial features extracted from the convolutional layer are subsequently passed through two LSTM layers with 256 neurons to learn the temporal structures of the input data. A recurrent dropout probability of 0.1 is also applied to each LSTM layer to prevent overfitting. Following this step, the output tensor from the second LSTM layer is flattened and directed into two fully connected dense layers. These layers aimed to obtain a higher-order representation of the processed features and enhance the separability of the instances into the different labels found within the training dataset. Both dense layers employ the rectified linear unit (ReLU) activation function and are initialized using the Glorot uniform weight initialization method, with the first dense layer having 64 neurons and the second layer containing 16 neurons.

The output of the second dense layer is finally fed into the final output layer with the same number of neurons as the number of user response labels (i.e., 3 neurons for thermal preference and air movement preference, and 2 neurons for thermal acceptability). The output layer utilizes a softmax activation function to yield a probability distribution of the target labels, where the predicted label y is determined by identifying the label with the highest predicted probability scores.

The training process of the base model architecture employs the Adam optimizer with a learning rate of 0.001 and utilizes the categorical cross-entropy loss function to calculate the model's prediction

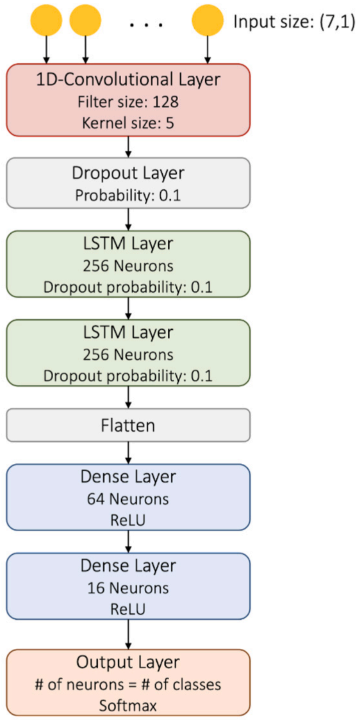

Fig. 3. Model architecture of the base CNNLSTM model with model hyperparameters.

error against the training data. The detailed architecture of the proposed CNNLSTM model is illustrated in Fig. 3.

3.5. Transfer learning model architecture

Utilizing the base CNNLSTM model architecture described in the previous subsection, two transductive transfer learning model architectures are developed for this study: transfer learning using fine-tuning (TL-CNNLSTM-FT) and transfer learning using feature extraction (TL-CNNLSTM-FE).

The training process of the TL-CNNLSTM-FT model begins by initializing a base CNNLSTM model, where the weights of each layer are randomly initialized using the Glorot uniform weight initialization method. This model is subsequently trained on the large amount of labeled data from the ASHRAE Global Thermal Comfort Database II (source domain), where the weights of the CNNLSTM layers are iteratively updated over 100 epochs through backpropagation. This process gradually reduces the model's prediction error in inferring the occupants' state of thermal comfort (i.e., thermal preferences, thermal acceptability, and air movement preferences), based on indoor and outdoor environmental conditions. Once the model has been pre-trained on the ASHRAE Global Thermal Comfort Database II (source domain), the TL-CNNLSTM-FT model is adapted to perform a similar task on the BCA Testbed Dataset (target domain) by retraining the model on small amounts of labeled data from the BCA dataset. During this retraining process, the weights of all layers are updated during the backpropagation process to fine-tune and improve the model's predictive performance in the target domain. Given that the predictive task between the source and target domain is identical, the final output layer of the pretrained TL-CNNLSTM-FT model was retained during the retraining process, instead of being replaced with new dense layers.

The TL-CNNLSTM-FE model follows a similar training process as the TL-CNNLSTM-FT model by first pre-training on the labeled data in the ASHRAE Global Thermal Comfort Database II (source domain), before retraining on the BCA Testbed Dataset (target domain). However, the key difference between both models lies in the retraining process, where instead of updating the weights for all model layers during the backpropagation step, only the weights of the final dense layers are updated, while the convolutional and LSTM layer weights are frozen. This retraining approach essentially allows the model to retain its feature-extracting abilities based on what was learned from the source domain.

Despite their differences, both models leverage the knowledge obtained from the information-rich source domain and adapt it to an information-poor target domain. This is achieved by improving the model's ability to extract informative high-level representations from the input features through parameter sharing in the pre-training step before enhancing the model's predictive performance in the target domain through retraining. Fig. 4 illustrates the model architectures of the TL-CNNLSTM-FT and TL-CNNLSTM-FE models, respectively.

3.6. Active transfer learning framework

To further enhance the transfer learning model's performance in both data efficiency and predictive accuracy, we introduced the concept of active learning during the model's retraining step to form the active transfer learning framework.

Specifically, within the model retraining step, an active learning algorithm is applied to the pool of unlabeled instances in the target domain to identify the most informative instances for human annotation and subsequent model retraining. The Query-By-Committee (QBC) sampling algorithm was adopted for this study due to findings from a prior study, which found that the QBC sampling algorithm was able to outperform other popular AL algorithms, such as Uncertainty Sampling, in reducing human annotation cost when applied to a similar dataset [24].

The QBC algorithm begins by defining a panel of classifiers, denoted as $C = \theta^{(1)}, ..., \theta^{(C)}$, which are trained on random samples of the labeled pool $P_{labeled}$. By applying this panel of classifiers to the unlabeled pool $P_{unlabeled}$, the unlabeled instance with the highest level of disagreement among the panel's predictions is considered to be the most informative instance and chosen for human labeling. The newly labeled instance is subsequently removed from $P_{unlabeled}$ and incorporated into $P_{labeled}$ before retraining the panel of classifiers on the updated $P_{labeled}$ for the next selection round. The vote entropy metric by [37] is employed in this study to quantify the degree of disagreement among the panel. The unlabeled instance X^* is chosen for human labeling based on the results of the following equation:

$$X^*_{QBC} = \underset{X}{\operatorname{argmax}} - \sum_{i}^{L} \frac{n(y_i)}{|C|} log \frac{n(y_i)}{|C|} \tag{2}$$

where $|C|$ represents the number of classifiers in the panel, L represents the number of class labels, and $n(y_i)$ denotes the number of classifiers categorizing instance X under class y_i.

Within our active transfer learning framework, we defined the ASHRAE Global Thermal Comfort Database II as our initial $P_{labeled}$, where a panel of classifiers was initially trained. Subsequently, we evaluated the unlabeled instances in the BCA Testbed Dataset, defined as our unlabeled pool $P_{unlabeled}$. By identifying the most informative instance in $P_{unlabeled}$ using the above equation, the instance was chosen for human labeling and included in $P_{labeled}$ for panel retraining and the next evaluation round.

The incorporation of active learning within the transfer learning framework enables us to focus the model retraining process on the most distinctive instances (i.e., most informative instances) in the target domain compared to the source domain, thereby enhancing the model's transferability in the new domain. Fig. 5 illustrates the model architecture of the proposed active transfer learning framework based on the TL-CNNLSTM-FT and TL-CNNLSTM-FE model architectures (ATL-CNNLSTM-FT and ATL-CNNLSTM-FE).

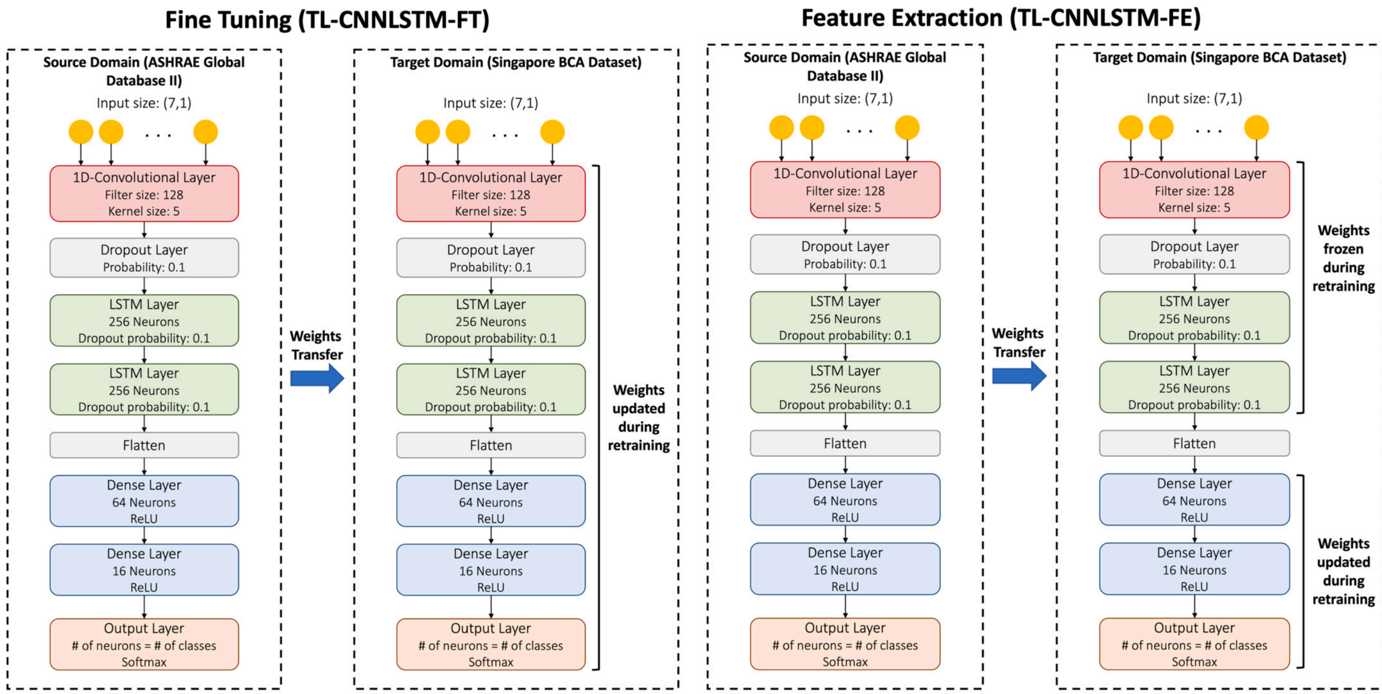

Fig. 4. Model architectures of the TL-CNNLSTM-FT and TL-CNNLSTM-FE with hyperparameters.

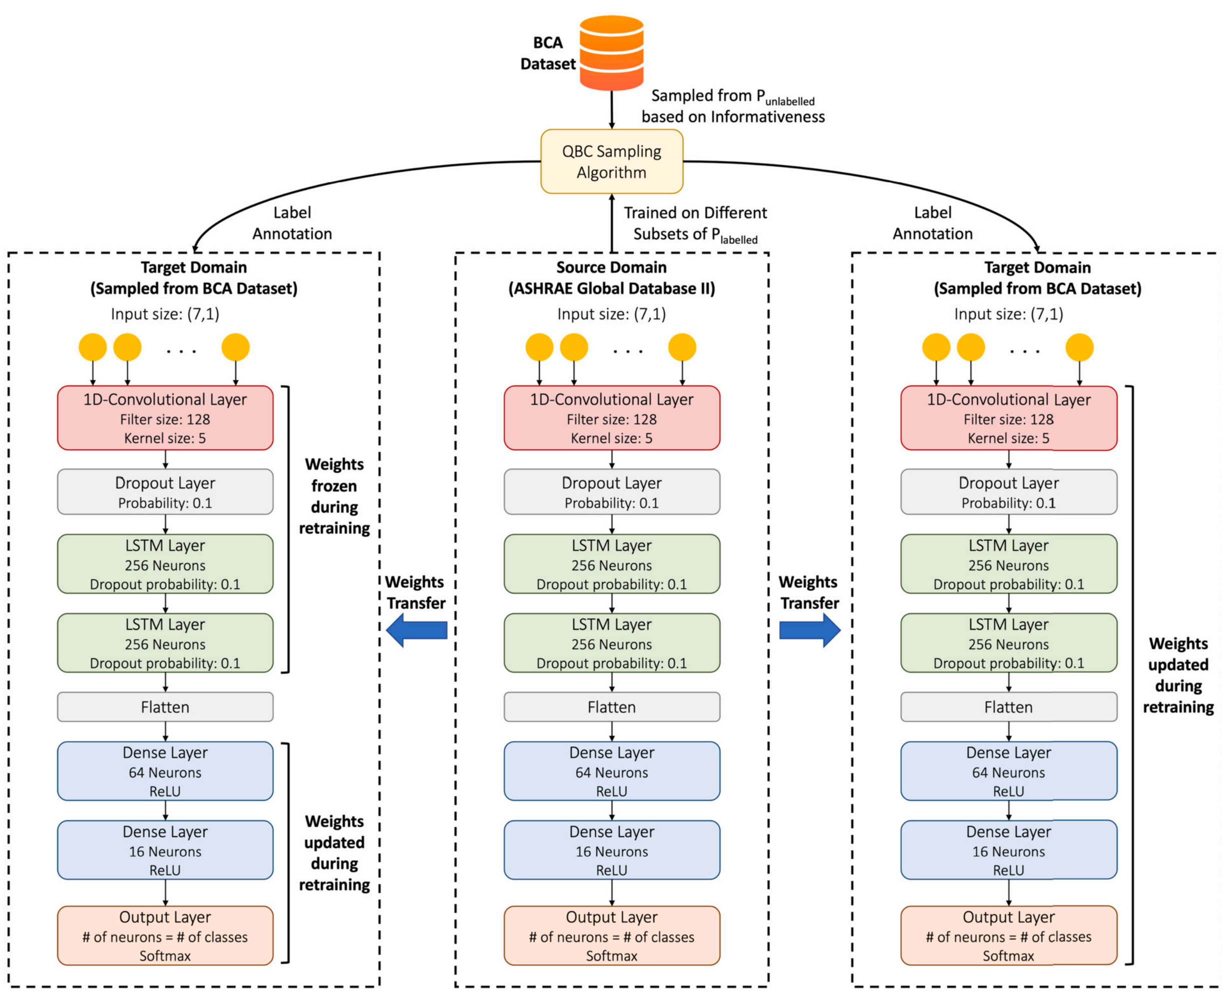

Fig. 5. Model architecture of the Active Transfer Learning Framework, including ATL-CNNLSTM-FT and ATL-CNNLSTM-FT with models hyperparameters.

Table 2
Details of model parameters used during implementation.

Model parameters	Value
Convolution Layer	Filter Size: 128 Kernel Size: 5
Dropout Layer	Dropout Probability: 0.1
LSTM Layer	Number of Neurons: 256 Recurrent Dropout Probability: 0.1
Dense Layer	Dense Layer 1: 64 Neurons Dense Layer 2: 16 Neurons Activation Function: ReLU Weight Initializer: Glorot uniform weight initialization
Output Layer	Number of Neurons: 2 (thermal acceptability) or 3 (thermal and air movement preference) depending on the classification task Activation Function: Softmax
Optimization Function	Adam Optimizer with learning rate of 0.001
Loss Function	Categorical Cross Entropy
Batch Size	128
Epochs	Pre-training Step: 100 Epochs Retraining Step: 30 Epochs with early stopping enabled and an evaluation window of 10 epochs

4. Results and discussion

4.1. Model implementation

The models outlined in this study were implemented using the Python programming language and the TensorFlow Keras library. Additionally, the hyperparameters chosen for model implementation are determined by experimenting with different hyperparameters on the model's predictive performance and findings from past studies that adopted the same model architecture. A summary of the model parameters and hyperparameters is described in Table 2.

During the model training process, the transfer learning models are pre-trained on 100% of the labeled data from the ASHRAE Global Thermal Comfort Database II. They are subsequently retrained on a limited number of labeled data from the BCA Testbed Dataset before being evaluated on a holdout set from the BCA dataset. The amount of labeled data used during the retraining step varies based on the experiments conducted, which are described in the following subsections. The holdout set is obtained by randomly setting aside 10 labeled instances from each study participant from the BCA testbed to ensure equal representation of each participant when evaluating the models' predictive performance.

4.2. Evaluation metrics

Two evaluation metrics were chosen for this study: accuracy and micro-average F1 score.

Accuracy is a popular classification evaluation metric used in past thermal comfort studies and is defined as the ratio of true positive instances for each label class against the total number of predicted instances across all label classes. The accuracy formula is formally expressed as:

$$Accuracy = \frac{\sum_{i=1}^{N} TruePositive_{ii}}{\sum_{i=1}^{N}\sum_{j=1}^{N} \ Predicted\ _{ij}} \tag{3}$$

In this equation, N represents the number of label classes, $TruePositive_{ii}$ is the number of correctly predicted instances for class i, and $Predicted_{ij}$ is the number of instances predicted as class j.

In addition, given the significant imbalance between the occupants' reported comfort labels in the BCA Testbed Dataset, as highlighted in Section 3.3, the micro-average F1 score was also adopted in this study to evaluate the comfort model's predictive performance across all label classes. As this evaluation metric is sensitive towards how well the model performs across all label classes, particularly for the minority class, it provides a more holistic assessment of the model's overall performance. By comparing this metric with the model's reported accuracy score, we could also determine if the model is biased towards the majority class based on the difference between both metrics.

The micro-average F1 score is calculated based on the harmonic mean of the precision $Precision_i$ and recall $Recall_i$ scores for each label class i as shown in the following equations.

$$\text{Precision}_i = \frac{\sum_{i=1}^{N} TruePositive_{ii}}{\sum_{i=1}^{N} TruePositive_{ii} + \sum_{j\neq i} FalsePositive_{ij}} \tag{4}$$

$$\text{Recall}_i = \frac{\sum_{i=1}^{N} TruePositive_{ii}}{\sum_{i=1}^{N} TruePositive_{ii} + \sum_{j\neq i} FalseNegative_{ij}} \tag{5}$$

$$\text{Micro-average F1} = \frac{2 \cdot \sum_{i=1}^{N} \text{Precision}_i \cdot \text{Recall}_i}{\sum_{i=1}^{N} (\text{Precision}_i + \text{Recall}_i)} \tag{6}$$

4.3. Impact of transfer learning approaches for personal comfort models

To evaluate the influence of transfer learning on the model's predictive performance, we compared a base CNNLSTM model against the TL-CNNLSTM-FT and TL-CNNLSTM-FE models in predicting the occupants' thermal preference, thermal acceptability, and air movement preferences, as depicted in Table 3. The primary distinction between the base CNNLSTM model and the TL-CNNLSTM-FT and TL-CNNLSTM-FE models lies in their training data, where the base model is trained solely on the BCA Testbed Dataset (target domain), while the latter models undergo pre-training on the ASHRAE database (source domain) before retraining on the BCA dataset.

Based on the empirical results presented in Table 3, both transfer learning models (TL-CNNLSTM-FT and TL-CNNLSTM-FE) consistently outperformed the base CNNLSTM model in terms of accuracy and micro-average F1 score. Specifically, the transfer learning models outperformed the base CNNLSTM model by 1.3%-3.4% in terms of accuracy and 2%-3.5% for micro-average F1 score when predicting the occupants' thermal preference, thermal acceptability, and air movement preference. These results clearly underscore the value of transfer learning, where the model was able to leverage insights from the information-rich ASHRAE database during the pre-training step and adapt to the information-poor BCA dataset during the retraining step, ultimately enhancing its predictive performance in the target domain. The implication of this finding is that practitioners can effectively use transfer learning to improve the performance of existing models trained under information-poor conditions without collecting additional data in the target building. Additionally, when comparing the model's predictive performance when using different transfer learning approaches, both TL-CNNLSTM-FE and TL-CNNLSTM-FT exhibited similar performances, with the TL-CNNLSTM-FT model performing marginally better in terms of model accuracy, while the TL-CNNLSTM-FE model performed consistently better when considering micro-average F1 score. Lastly, given that the reported accuracy and micro-average f1 scores reported in Table 3 are comparable for different comfort tasks and conditions, this result provides strong evidence that the models are also not biased towards the majority class.

4.4. Impact of active learning on transfer learning models' performance

This section investigates the impact of incorporating active learning into the transfer learning framework through two experiments:

Experiment 1 (Random selection, 9.8% of data): The first experiment involves randomly selecting two instances per participant from the target domain, which constitutes 9.8% of the total training data. Specifically, one instance was collected under AC (air-conditioning) and the other under NV (natural

Table 3
Model performance comparison of different Transfer learning models against the Base model (CNNLSTM) for thermal preference, thermal acceptability, and air movement preference. The reported accuracy and micro-average f1 scores are averaged across all respondents in the target domain.

Comfort Task	Model	Source	Target	% of Training Data from Target	Accuracy	Micro-average F1 Score
Thermal Acceptability	CNNLSTM (Base)	N.A.	BCA	100%	0.828	0.800
	TL-CNNLSTM-FE	ASHRAE			0.832	**0.818**
	TL-CNNLSTM-FT	ASHRAE			**0.839**	0.813
Thermal Preference	CNNLSTM (Base)	N.A.	BCA	100%	0.675	0.636
	TL-CNNLSTM-FE	ASHRAE			**0.698**	**0.658**
	TL-CNNLSTM-FT	ASHRAE			**0.693**	0.654
Air Movement Preference	CNNLSTM (Base)	N.A.	BCA	100%	0.619	0.602
	TL-CNNLSTM-FE	ASHRAE			0.630	**0.614**
	TL-CNNLSTM-FT	ASHRAE			**0.633**	0.593

Table 4
Model performance comparison with and without active learning for thermal preference, thermal acceptability, and air movement preference. The reported accuracy and micro-average f1 scores are averaged across all respondents in the target domain.

Comfort Task	Approach (Random vs Active Learning)	Model	Source	Target	% of Training Data from Target	Accuracy	Micro-average F1 Score
Thermal Acceptability	Random	CNNLSTM (Base Model)	N.A.	BCA	9.8%	0.782	0.790
		TL-CNNLSTM-FE	ASHRAE		9.8%	0.809	0.806
		TL-CNNLSTM-FT	ASHRAE		9.8%	0.809	0.777
	Active Learning	ACNNLSTM (Base Model)	N.A.	BCA	9.8%	0.789	0.759
		ATL-CNNLSTM-FE	ASHRAE		9.8%	**0.837**	**0.814**
		ATL-CNNLSTM-FT	ASHRAE		9.8%	0.823	0.804
Thermal Preference	Random	CNNLSTM (Base Model)	N.A.	BCA	9.8%	0.587	0.568
		TL-CNNLSTM-FE	ASHRAE		9.8%	0.630	0.604
		TL-CNNLSTM-FT	ASHRAE		9.8%	0.619	0.594
	Active Learning	ACNNLSTM (Base Model)	N.A.	BCA	9.8%	0.632	0.597
		ATL-CNNLSTM-FE	ASHRAE		9.8%	0.654	0.633
		ATL-CNNLSTM-FT	ASHRAE		9.8%	**0.668**	**0.636**
Air Movement Preference	Random	CNNLSTM (Base Model)	N.A.	BCA	9.8%	0.530	0.496
		TL-CNNLSTM-FE	ASHRAE		9.8%	0.582	0.563
		TL-CNNLSTM-FT	ASHRAE		9.8%	0.579	0.550
	Active Learning	ACNNLSTM (Base Model)	N.A.	BCA	9.8%	0.523	0.435
		ATL-CNNLSTM-FE	ASHRAE		9.8%	**0.616**	**0.574**
		ATL-CNNLSTM-FT	ASHRAE		9.8%	0.614	0.571

ventilation). The two instances were then utilized as training data for the retraining step.

Experiment 2 (Active learning, 9.8% of data): The second experiment involves using active learning (QBC sampling algorithm) to select the two most informative instances, also constituting 9.8% of the total training data. Similarly, one instance was collected under AC (air-conditioning) and the other under NV (natural ventilation). The two instances were then utilized as training data for the retraining step.

By pretraining the CNNLSTM model on the ASHRAE database and retraining it on the selected instances from the target domain, the results for both experiments are presented in Table 4. According to the empirical results presented in Table 4, the active transfer learning models (ATL-CNNLSTM-FT and ATL-CNNLSTM-FE) consistently outperformed their transfer-learning-only counterparts (TL-CNNLSTM-FT and TL-CNNLSTM-FE) both in terms of model accuracy and micro-average F1 score. More specifically, the active transfer learning models were able to outperform the transfer-learning-only models by 1.7%-7.9% in terms of accuracy and 1.0%-7.1% for micro-average F1 score when predicting the occupants' thermal preference, thermal acceptability, and air movement preference. Notably, the improvement in predictive performance after incorporating active learning is consistently observed not only in both transfer learning models (TL-CNNLSTM-FE and TL-CNNLSTM-FT) but also from the base CNNLSTM models without transfer learning, thereby demonstrating its generalizability to different model architectures. These findings provide sufficient evidence of the effectiveness of active learning in improving the predictive performance of the transfer learning models by optimizing the retraining process and selecting the most informative instances in the information-poor target domain for model retraining. In fact, the results show that by combining active learning with transfer learning, we were able to obtain a model performance comparable to another model trained on more than 10 times the existing training data, significantly improving the model's data efficiency. Lastly, given that the reported accuracy and micro-average f1 scores reported in Table 4 are comparable for different comfort tasks and conditions, this result provides strong evidence that the models are not biased towards the majority class.

4.5. Impact of different amounts of training data in target domain on predictive performance

After evaluating the impact of transfer learning and active learning on the model's predictive performance, this subsection aims to evaluate how the model performs with different amounts of training data in the target domain for the model retraining step. By evaluating the models' performance when using varying amounts of training data from the target domain (i.e., 0%, 9.8% ≈ 2 instances/data-points per participant - representing one instance for each AC and NV mode, 36% ≈ 8 instances per participant - representing one instance per condition out of 8 given in Table 1, and 100%), the results are generated and visualized in Fig. 6.

One of the observations from Fig. 6 indicates that the predictive performances of all models generally follow a positive trend as the models are retrained on an increasing amount of training data from the target

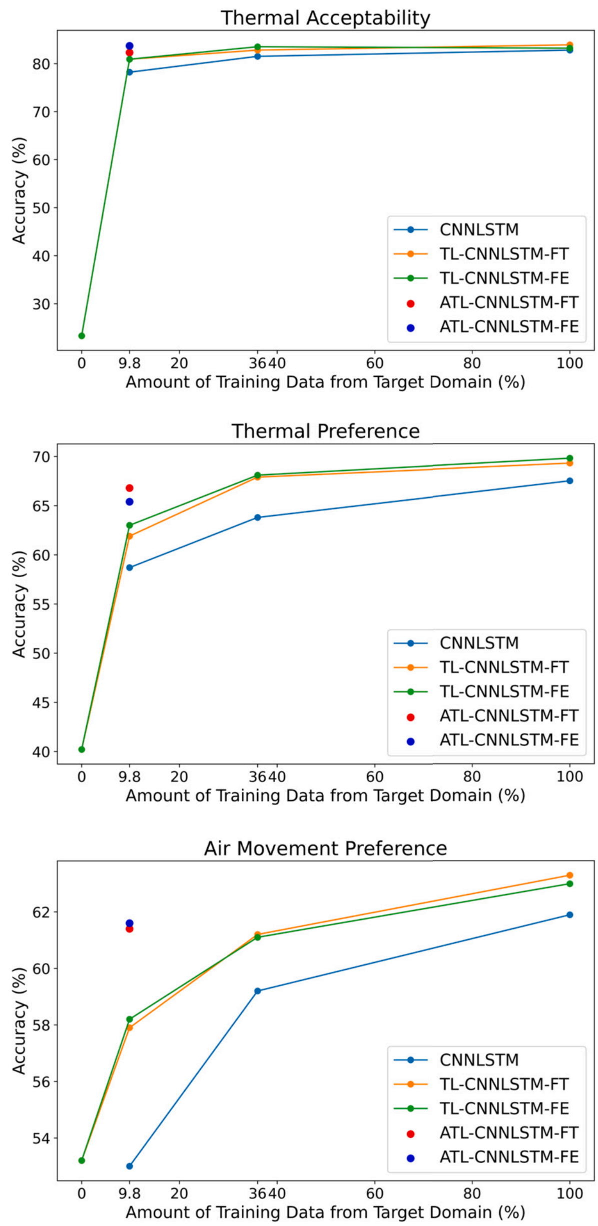

Fig. 6. Model performances of the base, transfer learning, and active transfer learning models after retraining on different amounts of data in the target domain for all comfort tasks. The reported accuracy scores on the Y-axis are averaged results across all respondents in the target domain.

domain. However, this trend does not follow a linear pattern but instead gradually converges, signifying diminishing returns in model improvement as more data from the target domain is available for retraining. This observation is likely due to the model having already captured the existing patterns within the dataset such that there are no additional patterns to learn from the new incoming data. This is especially true when it comes to modeling personal comfort as there is a physiological limit to the human body's ability to adapt to external environmental conditions. As such, the application of active learning is particularly effective in these cases as it helps to identify the most informative data points for learning these patterns. An implication of this result is that researchers conducting data collection studies to gather occupants' thermal comfort preferences should carefully balance their data collection budget with the marginal improvements from their models. This ensures optimal utilization of their limited budget to achieve greater gains in model performance.

Another notable observation arises from comparing the predictive performance of the base CNNLSTM model with that of the TL-CNNLSTM-FT and TL-CNNLSTM-FE models. It can be observed that the most significant improvement from the application of transfer learning occurs when very limited training data from the target domain is available for model retraining. This improvement gap gradually declines as more training data from the target domain becomes available. This result aligns with our intuition as the transfer-learned models were able to effectively bridge their knowledge gap in the information-poor target domain based on the knowledge gained in the source domain.

Finally, based on the performance boost resulting from the integration of active learning within the transfer learning framework, we can observe that ATL-based models with less than 10% of the training data were able to make predictions almost as accurately as the base CNNLSTM model trained with 100% of the data. In fact, in the case of thermal acceptability, ATL models with less than 10% of training data were observed in Fig. 6 to outperform traditional CNNLSTM models with 100% training data. Lastly, in terms of the ATL-based model's performance across different participants, the reported accuracy scores for all 58 respondents have a standard deviation of less than 0.16 across different comfort tasks, thereby demonstrating the personalized nature of the resulting model. This concludes that our proposed active transfer learning strategy was effective in improving the data efficiency of personal comfort models while maintaining adequate predictive performance.

5. Conclusion

We introduced a novel active transfer learning framework that integrates active learning and transfer learning to enhance the performance of personal comfort models in scenarios with limited data. The proposed framework was evaluated across three comfort tasks — thermal preference, thermal acceptability, and air movement preference — by leveraging knowledge from the extensive ASHRAE Global Thermal Comfort Database II and applying it to a smaller comfort dataset collected from the BCA testbed in Singapore.

Through a comprehensive evaluation of the comfort models developed using the proposed active transfer learning framework, the empirical results obtained led to multiple findings:

Firstly, the transfer-learning-enhanced models, namely TL-CNNLSTM-FT and TL-CNNLSTM-FE, were able to consistently outperform the base CNNLSTM model in terms of accuracy and micro-average F1 score. Specifically, these models exhibited improvements ranging from 1.3%-3.4% in accuracy and 2%-3.5% in micro-average F1 score when predicting the occupants' thermal preference, thermal acceptability, and air movement preference. Furthermore, when evaluating the models' performance using increasing amounts of training data, the most substantial improvement through transfer learning occurred when very limited training data from the target domain was available for model retraining, with diminishing returns observed as more data became accessible.

Secondly, the active transfer learning models, ATL-CNNLSTM-FT and ATL-CNNLSTM-FE, were able to further outperform their transfer-learning-only counterparts (TL-CNNLSTM-FT and TL-CNNLSTM-FE), achieving a boost of 1.7%-7.9% in accuracy and 1.0%-7.1% in micro-average F1 score. This result clearly demonstrated the effectiveness of combining active learning with transfer learning approaches in improving the occupants' personal comfort predictions.

Lastly, an intriguing finding emerged from comparing ATL-based models (ATL-CNNLSTM-FT and ATL-CNNLSTM-FE) and the base CNNLSTM model. The proposed active transfer learning models were able to make predictions almost as accurate as the base CNNLSTM model, using only less than 10% of the training data compared to the base CNNLSTM model trained with 100% of the data for all comfort tasks. This concludes that active transfer learning was an effective strategy for developing data-efficient personal comfort models using limited data while achieving sufficient predictive performance.

While our active transfer learning framework proves to be a valuable strategy for enhancing personal comfort models under limited data availability, it is also important to highlight some of the limitations of this study that will be addressed in future works. Firstly, given that the proposed active transfer learning framework is evaluated on a single target building in the tropics, there is a need to apply the framework across different target buildings from a range of climate zones to ensure the generalizability of the study's findings. Additionally, given that the adoption of active learning would inevitably increase the computation cost during the model training process, there is a need for practitioners to weigh the costs associated with data collection and the additional computation costs when deciding if the active transfer learning framework is appropriate for their use case. Lastly, other potential avenues for future research can also include devising more intelligent approaches for selecting relevant training data from the ASHRAE Global Database (i.e., source domain) during the initial training phase of transfer learning to further enhance the predictive performance of the resulting model when applied to the target domain.

To sum up, our active transfer learning framework proves to be a valuable strategy for enhancing personal comfort models, offering superior predictive performance even in scenarios with limited data availability. This study opens avenues for future research in data-efficient comfort modeling for practical applications.

CRediT authorship contribution statement

Zeynep Duygu Tekler: Writing – review & editing, Writing – original draft, Visualization, Methodology, Investigation, Conceptualization. **Yue Lei:** Writing – review & editing, Data curation. **Adrian Chong:** Writing – review & editing, Supervision, Resources, Project administration, Funding acquisition, Conceptualization.

Declaration of competing interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

Data availability

The research compendium for this article is available at https://github.com/ideas-lab-nus/ComfortPredict-ActiveTL, hosted on GitHub.

Acknowledgements

The National Research Foundation, Singapore, and the Ministry of National Development, Singapore, provided support for this research project through the Cities of Tomorrow R&D Programme (CoT Award COT-V4-2020-5). The views, opinions, findings, conclusions, or recommendations expressed in this material are solely those of the author(s) and do not reflect the opinions of these funding agencies.

References

[1] ASHRAE, Standard 55, Thermal Environmental Conditions for Human Occupancy, American Society of Heating, Ventilating, and Air Conditioning Engineers, Atlanta, Georgia, 2023.
[2] J. Kim, S. Schiavon, G. Brager, Personal comfort models–a new paradigm in thermal comfort for occupant-centric environmental control, Build. Environ. 132 (2018) 114–124.
[3] W. O'Brien, A. Wagner, M. Schweiker, A. Mahdavi, J. Day, M.B. Kjærgaard, S. Carlucci, B. Dong, F. Tahmasebi, D. Yan, et al., Introducing iea ebc annex 79: key challenges and opportunities in the field of occupant-centric building design and operation, Build. Environ. 178 (2020) 106738, https://doi.org/10.1016/j.buildenv.2020.106738.
[4] J. Xiong, A. Tzempelikos, I. Bilionis, P. Karava, A personalized daylighting control approach to dynamically optimize visual satisfaction and lighting energy use, Energy Build. 193 (2019) 111–126.
[5] Z.D. Tekler, R. Low, C. Yuen, L. Blessing, Plug-mate: an iot-based occupancy-driven plug load management system in smart buildings, Build. Environ. 223 (2022) 109472.
[6] Y. Lei, S. Zhan, E. Ono, Y. Peng, Z. Zhang, T. Hasama, A. Chong, A practical deep reinforcement learning framework for multivariate occupant-centric control in buildings, Appl. Energy 324 (2022) 119742, https://doi.org/10.1016/j.apenergy.2022.119742.
[7] D. Zhuang, V.J. Gan, Z.D. Tekler, A. Chong, S. Tian, X. Shi, Data-driven predictive control for smart hvac system in iot-integrated buildings with time-series forecasting and reinforcement learning, Appl. Energy 338 (2023) 120936.
[8] E. Ono, K. Mihara, K.P. Lam, A. Chong, The effects of a mismatch between thermal comfort modeling and hvac controls from an occupancy perspective, Build. Environ. 220 (2022) 109255.
[9] H. Zhang, A. Tzempelikos, X. Liu, S. Lee, F. Cappelletti, A. Gasparella, The impact of personal preference-based thermal control on energy use and thermal comfort: field implementation, Energy Build. 284 (2023) 112848.
[10] J. Kim, Y. Zhou, S. Schiavon, P. Raftery, G. Brager, Personal comfort models: predicting individuals' thermal preference using occupant heating and cooling behavior and machine learning, Build. Environ. 129 (2018) 96–106.
[11] F. Tartarini, S. Schiavon, M. Quintana, C. Miller, Personal comfort models based on a 6-month experiment using environmental parameters and data from wearables, Indoor Air 32 (11) (2022) e13160, https://doi.org/10.1111/ina.13160.
[12] M. Luo, J. Xie, Y. Yan, Z. Ke, P. Yu, Z. Wang, J. Zhang, Comparing machine learning algorithms in predicting thermal sensation using ashrae comfort database ii, Energy Build. 210 (2020) 109776.
[13] Q.Y. Li, J. Han, L. Lu, A random forest classification algorithm based personal thermal sensation model for personalized conditioning system in office buildings, Comput. J. 64 (3) (2021) 500–508.
[14] Y. Wu, J. Fan, B. Cao, Comparison among different modeling approaches for personalized thermal comfort prediction when using personal comfort systems, Energy Build. 285 (2023) 112873.
[15] Q. Chai, H. Wang, Y. Zhai, L. Yang, Using machine learning algorithms to predict occupants' thermal comfort in naturally ventilated residential buildings, Energy Build. 217 (2020) 109937.
[16] D. Li, C.C. Menassa, V.R. Kamat, Personalized human comfort in indoor building environments under diverse conditioning modes, Build. Environ. 126 (2017) 304–317.
[17] J.-H. Choi, D. Yeom, Study of data-driven thermal sensation prediction model as a function of local body skin temperatures in a built environment, Build. Environ. 121 (2017) 130–147.
[18] T.C. Cheung, S. Schiavon, E.T. Gall, M. Jin, W.W. Nazaroff, Longitudinal assessment of thermal and perceived air quality acceptability in relation to temperature, humidity, and co2 exposure in Singapore, Build. Environ. 115 (2017) 80–90.
[19] C. Shan, J. Hu, J. Wu, A. Zhang, G. Ding, L.X. Xu, Towards non-intrusive and high accuracy prediction of personal thermal comfort using a few sensitive physiological parameters, Energy Build. 207 (2020) 109594.
[20] Y. Feng, S. Liu, J. Wang, J. Yang, Y.-L. Jao, N. Wang, Data-driven personal thermal comfort prediction: a literature review, Renew. Sustain. Energy Rev. 161 (2022) 112357.
[21] P. Jayathissa, M. Quintana, M. Abdelrahman, C. Miller, Humans-as-a-sensor for buildings—intensive longitudinal indoor comfort models, Buildings 10 (10) (2020) 174.
[22] Y. Lei, Z.D. Tekler, S. Zhan, C. Miller, A. Chong, Experimental evaluation of thermal adaptation and transient thermal comfort in a tropical mixed-mode ventilation context, Build. Environ. 248 (2024) 111043.
[23] J. Ngarambe, G.Y. Yun, M. Santamouris, The use of artificial intelligence (ai) methods in the prediction of thermal comfort in buildings: energy implications of ai-based thermal comfort controls, Energy Build. 211 (2020) 109807.
[24] Z.D. Tekler, Y. Lei, Y. Peng, C. Miller, A. Chong, A hybrid active learning framework for personal thermal comfort models, Build. Environ. 234 (2023) 110148.
[25] Z.D. Tekler, A. Chong, Occupancy prediction using deep learning approaches across multiple space types: a minimum sensing strategy, Build. Environ. 226 (2022) 109689.
[26] Z.D. Tekler, Y. Lei, X. Dai, A. Chong, Enhancing Personalised Thermal Comfort Models with Active Learning for Improved hvac Controls, Journal of Physics: Conference Series, vol. 2600, IOP Publishing, 2023, p. 132004.
[27] N. Gao, W. Shao, M.S. Rahaman, J. Zhai, K. David, F.D. Salim, Transfer learning for thermal comfort prediction in multiple cities, Build. Environ. 195 (2021) 107725.
[28] N. Somu, A. Sriram, A. Kowli, K. Ramamritham, A hybrid deep transfer learning strategy for thermal comfort prediction in buildings, Build. Environ. 204 (2021) 108133.
[29] H. Park, D.Y. Park, Prediction of individual thermal comfort based on ensemble transfer learning method using wearable and environmental sensors, Build. Environ. 207 (2022) 108492.
[30] K. Li, Y. Liu, L. Chen, W. Xue, Data efficient indoor thermal comfort prediction using instance based transfer learning method, Energy Build. (2024) 113920.
[31] A. Natarajan, E. Laftchiev, A transfer active learning framework to predict thermal comfort, Int. J. Progn. Health Manag. 10 (3) (2019).
[32] V.F. Ličina, T. Cheung, H. Zhang, R. De Dear, T. Parkinson, E. Arens, C. Chun, S. Schiavon, M. Luo, G. Brager, et al., Development of the ashrae global thermal comfort database ii, Build. Environ. 142 (2018) 502–512.

[33] R. Low, Z.D. Tekler, L. Cheah, Predicting commercial vehicle parking duration using generative adversarial multiple imputation networks, Transp. Res. Rec. 2674 (9) (2020) 820–831.

[34] Z.D. Tekler, E. Ono, Y. Peng, S. Zhan, B. Lasternas, A. Chong, Robod, Room-Level Occupancy and Building Operation Dataset, Building Simulation, vol. 15, Springer, 2022, pp. 2127–2137.

[35] N.V. Chawla, K.W. Bowyer, L.O. Hall, W.P. Kegelmeyer, Smote: synthetic minority over-sampling technique, J. Artif. Intell. Res. 16 (2002) 321–357, https://doi.org/10.1613/jair.953.

[36] R. Low, L. Cheah, L. You, Commercial vehicle activity prediction with imbalanced class distribution using a hybrid sampling and gradient boosting approach, IEEE Trans. Intell. Transp. Syst. 22 (3) (2020) 1401–1410.

[37] I. Dagan, S.P. Engelson, Committee-based sampling for training probabilistic classifiers, in: Machine Learning Proceedings 1995, Elsevier, 1995, pp. 150–157.