

A Transfer Active Learning Framework to Predict Thermal Comfort

Annamalai Natarajan^{1*}, Emil Laftchiev²

¹ Philips Research, Cambridge, MA, 02141, USA
Annamalai.Natarajan@philips.com

² Mitsubishi Electric Research Labs, Cambridge, MA, 02139, USA
Laftchiev@merl.com

ABSTRACT

Personal thermal comfort is the feeling that individuals have about how hot, cold or comfortable they are. Studies have shown that thermal comfort is a key component of human performance in the work place and that personalized thermal comfort models can be learned from user labeled data that is collected from wearable devices and room sensors. These personalized thermal comfort models can then be used to optimize the thermal comfort of room occupants to maximize their performance. Unfortunately, personalized thermal comfort models can only be learned after extensive dataset collection and user labeling. This paper addresses this challenge by proposing a transfer active learning framework for thermal comfort prediction that reduces the burdensome task of collecting large labeled datasets for each new user. The framework leverages domain knowledge from prior users and an active learning strategy for new users that reduces the necessary size of the labeled dataset. When tested on a real dataset collected from five users, this framework achieves a 70% reduction in the required size of the labeled dataset as compared to the fully supervised learning approach. Specifically, the framework achieves a mean error of 0.82 ± 0.05 , while the supervised learning approach achieves a mean error of 0.85 ± 0.04 .

1. INTRODUCTION AND BACKGROUND

Personal thermal comfort is feeling that individuals have about how hot, cold or comfortable they are. Importantly, thermal comfort is a good predictor of human performance. Studies have shown that making office workers comfortable is critical to improving worker productivity and improving the office environment. In one study, Hedge et. al. (Hedge, Wafa, & Anshu, 2005) showed that reducing temperatures such that the average female office worker felt chilly increased her typing mistakes

*Research performed while at Mitsubishi Electric Research Labs. Annamalai Natarajan et al. This is an open-access article distributed under the terms of the Creative Commons Attribution 3.0 United States License, which permits unrestricted use, distribution, and reproduction in any medium, provided the original author and source are credited.

by 74%. The same chilly office worker had a reduced output (productivity) of 46%. Another study by IJzerman and Semin (IJzerman & Semin, 2009) showed that warmth in the office environment encouraged closeness and friendliness. Thus, improving and maintaining thermal comfort in the office environment can yield significant benefits in terms of improving worker performance. In fact, the improved work performance was estimated to be as much a 12.5% increase in worker wages (Hedge et al., 2005).

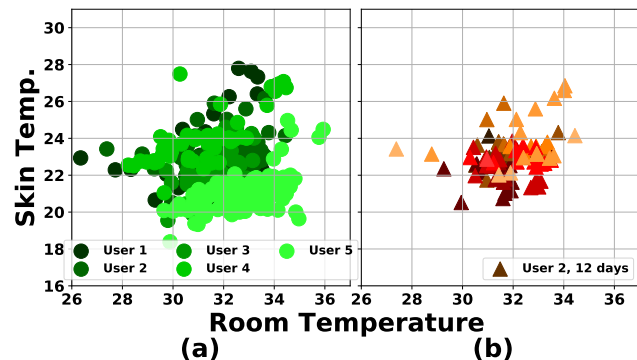


Figure 1. Scatter plot of skin temperature and room temperature (a) at the instant in time when five users reported being thermally comfortable (b) at the instant in time for one user over multiple days who reported being thermally comfortable

There are two principal difficulties when modeling thermal comfort. First, personal thermal comfort varies from one individual to the next. Often this variation can be explained by gender, ethnicity, location, and season. Second, personal thermal comfort varies within the individual because of their physical state, including conditions such as tiredness and sickness. Because of these difficulties, state of the art methods of thermal comfort estimation provide only coarse estimation of thermal comfort for large groups of individuals. Additionally, these models rely on a-priori assumptions about the composition of the group of occupants which increases the model error rates in today's diverse workforce (Belluck, 2015).

The variation in personal thermal comfort is illustrated in Figure 1. In subplot (a) five users' skin temperature and room temperature are shown, color coded, at instants of feedback. Here, feedback refers to thermal comfort ratings as provided by users. For these examples, the users reported being thermally comfortable. In this figure, one can clearly see the difference in preference among the five users. In subplot (b) the same plot is repeated for a single user's data across multiple days. Here one can clearly observe the variation in preference of a single user. In both cases, we note that the preferences overlap, underlining the fact that some part of the individual preferences is shared and necessarily common through the underlying human thermo-regulation, but preferences vary both between users and between days for the same subject.

Given the importance of thermal comfort in human performance, it is desirable to find an approach to overcome the difficulty in modeling so that we can understand and maximize personal thermal comfort. This requires the development of personalized thermal comfort models which rely on the availability of large quantities of labeled data examples. Such an approach was recently used by several authors (Laftchiev & Nikovski, 2016; Ranjan & Scott, 2016; Huang, Yang, & Newman, 2015) who all note the difficulty in securing the cooperation of users during the experiments. Thus obtaining labeled data examples is emerging as a key obstacle.

This paper addresses the problem of modeling thermal comfort with a minimal number of labeled examples. Our goal is to develop a machine learning framework that uses data collected from an IoT system (Laftchiev & Nikovski, 2016) to create accurate models that can be used to predict personalized thermal comfort. To achieve this goal a novel transfer active learning framework is presented. The first part of the framework leverages knowledge from a few base users (a group of initial users as part of a controlled experiment) using transfer learning to learn a general model of thermal comfort. This model is then personalized using data from a new user which is obtained through queries issued by an active learning algorithm. The results in this paper show that while still in the supervised machine learning setting, this approach greatly reduces the number of labeled examples that must be provided by a user.

The active learning research presented in this paper is set in the pool-based setting. In pool-based active learning, the assumption is that we have access to all unlabeled examples but only few examples are chosen to be queried for labels. This is the standard setting for the active learning problem. However, within this setting and within this paper's research careful attention is paid to choose specific techniques (and to tune those techniques) such that the transition to future work in the stream-based setting can be facilitated. Specifically, the methods chosen consider that in the stream-based setting user data is arriving continuously (as possible data examples that can be labeled by users), and the labeling window is inherently

short. However, a study in the stream-based setting remains to be performed at a future time. In this work, we evaluate the feasibility of using active learning techniques to minimize the labeling effort in thermal comfort prediction using wearable devices. To the best of our knowledge ours is the first work to demonstrate this in a user cohort.

The work presented in this paper is tested on a dataset collected from five users who were exposed to a variety of conditions from high temperature and high humidity to low temperature and low humidity. The results of this study show that it is possible to reduce the number of required labeled examples by 70%, on average over all five users. Specifically, our framework achieves a mean root mean square error (RMSE) over five users of 0.818 with standard error of ± 0.05 with just 25 labeled examples in comparison to a mean RMSE of 0.845 with standard error of ± 0.04 when using 82 labeled examples. This indicates that our framework can achieve the similar performance albeit with less resources.

In summary, this paper makes three contributions:

- A transfer active learning framework for thermal comfort prediction which leverages domain knowledge via transfer learning and minimizes the number of labeled examples via active learning.
- Two query-by-committee querying strategies for active learning in regression settings with a novel disagreement score.
- An empirical study using five participants which shows: (a) The feasibility of accurately modeling thermal comfort using machine learning for multiple individuals, (b) The feasibility of reducing the labeling effort of new participants using the described framework.

2. BACKGROUND AND RELATED WORK

This section begins by first providing some background on thermal comfort prediction. This is followed by a short background on transfer learning and active learning and concludes with a comparison between the method developed in this paper and previously published approaches. Because the volume of work addressed here is large, and our space is limited, we invite the reader to study transfer learning and active learning in more depth in (Pan & Yang, 2010) and (Settles, 2010) respectively.

2.1. Thermal Comfort Estimation

In this paper our goal is to build on prior work in (Laftchiev & Nikovski, 2016; Ranjan & Scott, 2016; Huang et al., 2015) to model individual personal thermal comfort. The first study to model thermal comfort was performed by Povl Ole Fanger (Fanger, 1967; *Ergonomics of the thermal environment – Analytical determination and interpretation of thermal comfort using calculation of the PMV and PPD indices and local*

thermal comfort criteria, 2005). Fanger did not focus on the nuance of modeling a single individual and instead modeled the mean thermal comfort vote of a group of people. During Fanger's experiment, thermal comfort feedback was given on an integer scale between 1 and 7. This scale is typically called the Bedford Scale or when offset to be between -3 and 3, the ASHRAE scale. Fanger's model is calibrated such that at most 5% of respondents are dissatisfied when the predicted mean vote is comfortable. This model, developed in the 1970s, was later adopted as an international standard in ISO 7730.

Fanger's model is based on heat balance equations that describe the transfer of heat from the body to the environment, with model constants learned from a group study. This model requires one input, room temperature, and makes assumptions about other input factors such as metabolic rate, effective mechanical power produced by the body, clothing insulation, surface area of the body, mean radiant temperature, relative air velocity, humidity, convective heat transfer, and clothing surface temperature. These assumptions are also at the core of the critiques of the model since they were made based on a very homogeneous group of individuals.

Another class of models called adaptive models that explain thermal comfort as a function of outdoor and indoor temperature have been proposed. Examples of this literature include the European Committee for Standardizations CEN method (*Indoor Environmental Input Parameters for Design and Assessment of Energy Performance of Buildings - Addressing Indoor Air Quality, Thermal Environment, Lighting and Acoustics*, 2006), and the American Society of Heating, Refrigerating, and Air-Conditioning Engineers (ASHRAE) method (*Thermal Environmental Conditions for Human Occupancy*, 2013). Building on these approaches, Haldi proposed a probabilistic model for thermal comfort (Haldi, 2010). These models are typically calibrated by season (*Summer, Winter, etc.*) and address the critique of the physiological models that they are incapable of capturing seasonal variation in individual preferences. Lastly, data driven approaches have been proposed by Jiang and Yao (Jiang & Yao, 2016) and Farhan et al. (Farhan, Pattipati, Wang, & Luh, 2015), which these focus on a few machine learning models or the prediction of comfort on a limited scale, respectively.

All modeling efforts suffer from the drawback of necessarily requiring assumptions about the individual. However, with the advent of wearable sensors and Internet of Things (IoT) technology, it is now possible to measure many of the variables estimated in the prior models. Recent work by Laftchiev and Nikovski (Laftchiev & Nikovski, 2016) seeks to capture this possibility by designing a new IoT system to explicitly sense as many features as possible from an individual user and then use supervised machine learning to identify an individual's model of thermal comfort. This work showed that using biometric measurements from wearable devices and am-

bient measurements of temperature, humidity, and airspeed, and feedback from the user, it is possible to accurately model individual thermal comfort. The caveat is that the user must provide a sufficiently large number of labeled examples in order to develop accurate prediction models.

Another study that employed wearable sensors and room sensors was performed by Huang et al. (Huang et al., 2015). In this study the authors did not use the standard scale of comfort measurement (the ASHRAE or Bedford Scale). Instead they proposed a new five level scale that combined thermal and comfort sensation indices. A classifier is then trained on features extracted from the sensors data in duration intervals of 5 and 30 minutes. Notably, these authors also point out that the main challenge in developing personalized thermal comfort prediction models for a user is the lack of user labeling across a diverse number of conditions.

A third study of interest was performed by Ranjan et al. (Ranjan & Scott, 2016) using thermographic images for thermal comfort prediction. Here the authors leverage the intuition that blood rushes in to parts of the body to keep warm and rushes out to keep the body cool. Thus heatmaps of the body can be constructed from the thermographic images to reveal whether a person is hot or cold. During the course of this study, participants had thermographic images taken twice per day in their office. At the same time, each participant provided labeled examples about their comfort. Despite the small labeling effort required, the authors noted the difficulty in obtaining labeled examples from users - the minimum labeled examples per user collected were 8, the median number of labeled examples collected was 24, while the maximum number collected was 33. The expected number of labeled examples per user was 50.

2.2. Transfer Learning for Regression

Transfer learning is a type of machine learning where knowledge from one domain is transferred to another with the goal of facilitating learning. In our problem setup, given N users, domains refer to different users. Specifically the source domain pertains to data from $N - 1$ users and the target domain refers to data from the N^{th} user. Predicting thermal comfort falls under inductive transfer learning where labeled data are available in both source and target domains but the crucial difference is we do not assume we have access to all labeled data in target domain.

One of approaches to inductive transfer learning is parameter transfer - where the assumption is that parameters for individual prediction models for similar tasks should be sampled from the same prior distribution (Pan & Yang, 2010). Parameter transfer has been shown to work for classification models like support vector machines (Evgeniou & Pontil, 2004) and conditional random field models (Natarajan et al., 2014).

For regression problems, parameter transfer has been restricted to Gaussian Processes (GP) (Bonilla, Chai, & Williams, 2008; Schwaighofer, Tresp, & Yu, 2005). The general idea is to share a GP prior that captures the dependencies between different tasks and/or domains. This approach is ideal when representative data is simultaneously available to jointly learn a GP prior. This is challenging in the present problem setup because it is assumed data in the target domain is scarce.

Hence our approach to parameter sharing is sequential where we first learn the parameters in the source domain and utilize this information as data becomes available in the target domain. Specifically, our approach to parameter sharing is to first learn the source domain parameters and second to penalize the deviation of target domain model parameters from source domain model parameters. This has the added advantage that in the absence of target domain data the prediction model would fall back on the source domain model to make predictions. This is clearly better than making predictions using a model initialized with random parameters, or assuming a baseline comfort score.

2.3. Active Learning for Regression

Active learning is a type of machine learning where a prediction model achieves good performance when it is allowed to choose the examples from which to learn (Lewis & Catlett, 1994). An active learner chooses a sample to be labeled via querying and then requests an oracle to provide a label for the chosen sample. The majority of published results in active learning focus on classification problems, in contrast, few papers address the work of developing active learning approaches for regression problems (Cai, Zhang, & Zhou, 2013).

Active learning for regression can be subdivided into model-free and model-based approaches. The model-free strategies are active learning approaches that do not rely on a prediction model to determine which data samples to label. Instead these approaches rely only on the statistics of the data distribution (O'Neill, Delany, & MacNamee, 2017). The most popular model-free approach is a density-based querying approach which seeks labeled examples for data points residing in high density regions of the dataset. These regions are hypothesized to be representative of the underlying data distribution and thus of labeling interest (Settles, 2010). The difficulty faced in model-free active learning approaches is that successive queries do not account for prior knowledge gained and often end up issuing redundant queries. Therefore the model-free active learning approach is not suitable in this paper because when the problem setting involves human user labeling there is an extreme constraint with respect to the number of queries that a user is willing to label.

On the other hand, most model-based active learning for the case of building a regression model builds on the early work by Geman et al. (Geman, Bienenstock, & Doursat, 2008), where

the generalization error (*e.g.*, *RMSE*, *MSE*, *etc*) is decomposed of three terms – error from model misspecification, error from labeling noise and error from model variance. The first two terms are fixed by the choice of prediction model and experiment design. For this reason, most of the published research has focused on minimizing the model variance such that the total generalization error (made up of all three components) is minimized. Typically model variance reduction techniques have relied on computing a variant of Fisher information which sets a lower bound on the variance of model parameter's estimates (Settles, 2010).

The challenge in using variance reduction techniques for regression is that the statistics such as Fisher information must be computed on the whole data distribution and is therefore not feasible to be computed when samples arrive sequentially. The last caveat is important to this work because the final goal of the framework presented in this paper is to be transplanted into the stream-based setting where knowledge of the complete data distribution is unknown.

A second approach to model-based active learning is query-by-committee (QBC). The goal of QBC is to minimize the space of possible predictions (also known as hypothesis, version space) given the current labeled dataset (Burbidge, Rowland, & King, 2007). To achieve this goal QBC relies on a committee to vote on available pool of examples and the most controversial example is chosen to be labeled. Once this sample is chosen, the committee members who disagreed the most with the provided ground truth label update their prediction models to minimize disagreement on similar data points in the future. Classically proposed QBC is unlikely to perform well in the case of thermal comfort model learning because the committee of users have fixed thermal comfort models which cannot simply be retrained when a new user provides a label.

2.4. Transfer and Active Learning Combined

The work most closely related to the active transfer learning framework proposed in this paper is that of Wang et al., (Wang, Huang, & Schneider, 2014). In this paper the authors develop a unified framework to perform active transfer learning for the case where a regression model is to be identified. The proposed approach is a domain adaptation approach that is based on the differences in both the marginal ($P(X)$) and conditional distributions ($P(Y|X)$) in source and target domains. To account for the difference in the marginal distribution, the authors perform a covariate shift and to account for the differences in the conditional distribution the authors propose two approaches both subject to the smoothness assumption: the first, is to match the conditional distributions between the source and target domains; the second, is to use the Gaussian Processes to model the source, target tasks and the offset in between. The authors leverage active learning to choose which examples in target domain should be labeled.

A second paper of importance where domain adaptation is performed and then is augmented by active learning is the work of Sugiyama et al., (Sugiyama, 2006). In this paper the authors assume that the marginal distributions differ but the conditional distributions remain the same between the source and target domains. The difference in the marginal distributions is viewed as a covariate shift problem and is addressed by computing importance weights such that the distribution is rebalanced in the source domain to match the distribution in the target domain. The authors then perform active learning, choosing examples to be labeled in the target domain, to learn the importance weights used in the distribution rebalancing problem. Concretely, the active learning approach is batch active learning where a subset of available examples is sampled to be labeled in a single shot. In practice, multiple subsets need to be chosen in order to identify the subset which minimizes generalization error.

There are two crucial differences between our work and these two prior work: first, both papers rely on computing importance weights to handle covariate shift. These importance weights are typically computed by estimating the probability densities of the marginal distribution in the source and target domain. This is generally challenging in datasets which are high dimensional but have a low sample counts. This is the case for the dataset in this paper. Second, both approaches operate in a pool-based settings where active learner can choose examples in small batches (Sugiyama et al., (Sugiyama, 2006)). Wang et al., (Wang et al., 2014) take advantage of the pool-based setting to identify regions of high utility from low utility and strategically place queries. Such approaches cannot be deployed to the stream-based settings and must be necessarily re-worked. In contrast, the work in this paper presents an approach that can be adapted the stream-based setting in a future study.

3. TRANSFER ACTIVE LEARNING FRAMEWORK FOR THERMAL COMFORT PREDICTION

Having reviewed the case for thermal comfort prediction as well as the relevant background in transfer learning and active learning, here we return to the problem of developing personalized thermal comfort model with a minimal number of labeled examples. This section develops the framework of transfer active learning that was described in the introduction. The overall goal is to leverage prior knowledge to rapidly learn a prediction model of thermal comfort and tune this model given only a few labeled examples of data.

3.1. Notation

To begin we present notation for the development of the framework. In this work we assume that we are given a dataset D which contains n labeled samples of the form $D = (x^i, y^i) \forall i \in \{1 \dots n\}$. Here each x^i corresponds to a feature vector,

each real valued, $x^i \in \mathbb{R}^k$, and each corresponding to data from wearable and ambient room sensors. The index i denotes the sample number while k denotes the length of the vector which corresponds to the number of features used in the prediction model. The target values y^i are drawn from a pre-defined set, $y^i \in \{0, \pm 1, \pm 2, \pm 3\}$. These correspond to thermal comfort rating given as feedback from the users.

The goal of this paper and the framework is to learn a prediction model, $\mathbf{h}, \mathbf{h} : x^i \rightarrow y^i$ that for any input vector x^i outputs a prediction target value $\hat{y}^i = \mathbf{h}(x^i)$. Because in this paper the prediction model will be learned using a regression method (more details under framework development), we stipulate that the predicted target value must not deviate more than ϵ , in the squared sense, from the actual target value y^i as $(y^i - \hat{y}^i)^2 < \epsilon \forall i$. For convenience the n labeled data samples are all expressed as matrix, which we call the design matrix, X , with n rows and k columns. We also represent all the labels as a vector y of length n .

3.2. Transfer Learning for Linear Regression

For the development in this work, the target values, y , are treated as continuous values that are restricted to the range $\{-3, +3\}$. The inherent assumption here is that while users are forced to discretize their state into 7 levels, in practice their thermal comfort is much more nuanced.

Furthermore, treating the problem of thermal comfort prediction as a regression problem addresses the problem of class imbalance. In particular because most users are in an HVAC controlled space, we anticipate that most feedback received will be in the range $\{-1, +1\}$ leading to severe class imbalances for the $-3, -2, +2, \text{ and } +3$ classes. Thus using regression methods is a natural approach when training thermal comfort predictors.

To provide maximum transparency to the model, this paper focuses on the problem of linear regression. Linear regression provides an easier quantification of the effect of each feature on the model output. Linear regression is parametrized by weight vector, W , such that the product of design matrix and weight vector results in \hat{y} as,

$$X^T W = \hat{y}. \quad (1)$$

The standard approach to finding the regressor weight vector is called ordinary least squares (OLS), where the goal of OLS is to minimize the squared sum of the differences between the estimated target values and the real target values. These differences are called the residuals and the sum of the residuals, often written as an optimization objective is expressed as,

$$\|X^T W - y\|_2^2. \quad (2)$$

The OLS estimate of W is prone to high variance in the model

weights and poor allocation (selection) of the weights among the features. Furthermore the classical, analytical solution to this problem is not well posed, suffering from numerical issues in the event that the design matrix is not easily invertible.

To remedy these issues, a penalty is introduced on the regressor weight vector. In this paper we choose this penalty to take the form of the ℓ_2 -norm, which means that the derivation below follows the Ridge Regression framework. Here the ℓ_2 -norm is chosen because of its more beneficial treatment of correlated features (Hastie, Tibshirani, & Friedman, 2009). The added penalty parameter reduces the model variance and results in a solution where some feature weight may be close to zero. This minimizes overfitting and reduces variance in estimates. The new objective function to solve is thus,

$$\|X^T W - y\|_2^2 + \lambda \|W\|_2^2. \quad (3)$$

In equation (3), λ , is the penalty parameter that determines the weight of the penalty term in the solution. Increasing λ leads to smaller weight coefficients in W , and decreasing λ leads to larger weight coefficients in W . Because of this, λ is said to control the shrinkage of the regressor coefficients.

Classically, when utilizing Ridge Regression, the shrinkage parameter is optimized such that the coefficients are driven towards zero without compromising the model error performance. This approach to Ridge Regression has a Bayesian interpretation where the weight vector coefficients are sampled from a prior normal distribution with mean zero and variance = $\sqrt{\frac{1}{\lambda}}$.

An alternate approach to Ridge Regression is to shrink the coefficients towards a non-zero prior distribution. When this approach is taken, the non-zero prior distribution represents some prior knowledge about the problem. In this case, it is said that the shrinkage of the coefficients toward the prior distribution induces a transfer of domain knowledge because the weight vector we find should be as close to the prior distribution as possible. The modified ridge regression problem has the following form,

$$\|X^T W - y\|_2^2 + \lambda \|W - W_p\|_2^2. \quad (4)$$

In equation (4), W_p is a vector containing a sample regressor vector. This vector represents the mean of the prior distribution described above. Note that setting W_p to zeros results in the classical ridge OLS from equation (3).

Multiple approaches exist to estimate the prior regressor, W_p . In this paper, we posit that when the goal is to derive a personalized thermal comfort model, we can assume that there are strong similarities between users, and that the model must only be slightly modified to fit a new individual. This assumption is rooted in the physiology of thermoregulation,

which does not differ from one person to the next. It is simply the preferences of the individual that differ.

One convenient prior for transfer learning in the case of thermal comfort modeling is a general thermal comfort over a group of users. That is, suppose that we have N datasets collected from N distinct users. Then we can find a general linear regressor, using equation (3), that describes the data from $N - 1$ users. We call this regressor our population model, W_p . We then use equation (4).

Solving equation (4), will then yield the personalized thermal comfort model for the N^{th} user. This approach to introducing a prior intuitively captures the idea that new user's coefficients, W , should be very similar to other users while allowing for individual differences.

Setting this up as an optimization problem, the ridge regression coefficients are learned by minimizing the following objective function,

$$\hat{W}^{ridge} = \arg \min_W (\|X^T W - y\|_2^2 + \lambda \|W - W_p\|_2^2). \quad (5)$$

In this formulation, the first term is the loss function, which has the usual format of equation (3), the second term penalizes the deviation of ridge coefficients of the new model W from the prior model W_p . Taking the derivative of this objective with respect to the new regressor weight vector W and setting it equal to zero results in analytical solution, which we term modified ridge regression,

$$\hat{W}^{ridge} = (X^T X + \lambda I)^{-1} (X^T Y + \lambda W_p). \quad (6)$$

3.3. Incorporating Active Learning

The goal of this framework is to create regression models that predict personal thermal comfort but do not require large quantities of labeled examples per user. So far we have introduced the transfer learning component of the framework, however, in order to personalize the model to the N^{th} user, this user must provide some feedback. Combining active learning with transfer learning is a logical approach to reducing the labeling effort for thermal comfort modeling.

In pool-based active learning, solutions often begin with the introduction of, \mathcal{A} , the pool of all available examples that are yet to be labeled and, \mathcal{L} , the set of labeled examples which are chosen through some active learning strategy. Importantly, in the pool-based setting all labels exist, but there is some cost associated to obtain a label that is to be minimized through sample selection. The overall goal of active learning is to choose an optimal subset of m (where $m \ll n$) labeled examples \mathcal{L} such that it achieves good generalization performance on the test set.

There are two important components of active learning; the

Algorithm 1 Transfer Active Learning

```

1: procedure THERMAL_COMFORT( $N, Budget$ )
2:    $Train \leftarrow \{\}$ 
3:    $Available \leftarrow Data(N)$ 
4:    $Other \leftarrow Data(N - 1)$ 
5:    $W_p \leftarrow Ridge(Other)$ 
6:    $Ridge_{W_p} \leftarrow random\ weights$ 
7:   while  $Budget \neq 0$  &  $Available \neq \emptyset$  do
8:      $(x, y) \leftarrow query\_strategy(Ridge_{W_p}, Available)$ 
9:      $Train \leftarrow Train \cup (x, y)$ 
10:     $Ridge_{W_p} \leftarrow Ridge_{W_p}(Train)$ 
11:     $Available \leftarrow Available - (x, y)$ 
12:     $Budget \leftarrow Budget - 1$ 
13:  end while
14:  return  $Ridge_{W_p}$ 
15: end procedure

```

\triangleright Predict Thermal comfort for user N using labeling budget $Budget$
 \triangleright Labeled examples \mathcal{L} of the N^{th} user
 \triangleright Available examples \mathcal{A} of the N^{th} user
 \triangleright Data from $N - 1$ users
 \triangleright W_p is learned on $N - 1$ users; Equation 3
 \triangleright Initial model; W_p is used from line #5 and used as in Equation 5
 \triangleright $Budget$ not empty and sample available to query
 \triangleright Pick and return labeled example using querying strategy
 \triangleright Add labeled example to \mathcal{L}
 \triangleright Update model; W_p is used from line #5 and used as in Equation 5
 \triangleright Remove labeled example from \mathcal{A}
 \triangleright Decrease budget
 \triangleright Return final model

labeling budget and the querying strategy. The labeling budget is simply the total number of labeled examples that can be obtained. In the context of personalized thermal comfort modeling this is the number of labeled examples that each user is allowed to be asked. Because in this problem the user should not be disturbed frequently, the labeling budget should be as small as possible.

The querying strategy is the approach used to determine which examples in the set \mathcal{A} should be labeled. In this paper we propose a modified QBC approach. In a typical QBC approach, the labeled dataset \mathcal{L} is used to update the committee members. Here we choose not to update the committee members, but instead we update only the N^{th} user's current predictive model. There are two reasons for choosing to update N^{th} user's predictive model: first, a labeled example from the N^{th} user could benefit only those committee members that exhibit a significant overlap in thermo-regulatory behavior. The consequence of using labeled examples to update committee members who are significantly different will result in noisy predictions when issuing subsequent queries; second, the goal of this work is to develop personalized prediction models with as few labeled examples as possible and hence updating the N^{th} user's predictive model gets us towards that goal quickly. The proposed QBC strategy is thus to choose examples which cause the committee members and the N^{th} user's predictive model to maximally disagree. Intuitively this means that the proposed QBC technique prefers examples for which the N^{th} user's model is uncertain about but the committee is fairly certain about.

A key point to address here is the notion of disagreement. We choose to define a disagreement score, d_i , for the i^{th} example in \mathcal{A} is computed as,

$$d_i = \left| \left(\frac{1}{C} \sum_{c=1}^C \hat{y}_c^i \right) - y_{\mathcal{L}}^i \right| \quad (7)$$

In equation (7), C is the number of committee members, \hat{y}_c^i is the prediction associated to the c^{th} committee member and $y_{\mathcal{L}}^i$ corresponds to the prediction made by the N^{th} user's prediction model which has been trained only using the labeled examples, \mathcal{L} , obtained thus far. We compute d_i for all examples in \mathcal{A} and pick the example with the maximum disagreement score to be labeled. It is important to note that the disagreement score defined in this manner prefers labeled examples that predict thermal comfort with opposite signs, for example -2 predicted as $+2$, versus examples that predict thermal comfort with the same sign but differ significantly in magnitude, for example -3 predicted as -2 . This disagreement score accommodates individual differences in thermo-regulatory behavior, for example the layering of clothes, while focusing on difference that may arise in dataset collected from different individuals; for example N^{th} user's model predicts cold when all other users feel hot under similar conditions.

Combining the transfer learning and the active learning, the complete transfer active learning framework is presented in Algorithm 1. In Algorithm 1, the function 'Thermal_Comfort' has two arguments: the user, N , for whom we personalize thermal comfort predictive model and labeling budget that determines the maximum number of active learning queries to be issued. The algorithm then creates, \mathcal{A} , a pool of available data examples from the N^{th} user and combines data from $N - 1$ users into 'Other' (lines #2 - #4). Then in line #5, the data from $N - 1$ users is used to train a ridge regression model which is the population level model, W_p . W_p is used then to learn the N^{th} user's initial model via transfer learning in line #6.

After the initial models are created, the algorithm iterates $Budget$ times, discovering a point to label in each iteration. To do this, a label is queried from the pool of available examples (line #8), the training dataset is updated (line #9) and the ridge regression model penalized by the population model is retrained (line #10).

The loop ends when the training budget is exhausted or there are no more samples available to label. While this loop is running, the algorithm updates the available dataset and adjust the budget of remaining labeled examples (lines #11 – #12). When the algorithm terminates, an updated ridge regression model is returned.

4. EMPIRICAL PROTOCOLS

Having developed the active transfer learning framework this section introduces the empirical protocols used to verify the approach. In particular, this section first introduces the participants which took part in the study. Then the experimental setup is discussed, including the sensors that were used and the feedback method which the users used. The data preprocessing is described along with data partitioning to train and test the framework. The section concludes by describing the experimental methods used for validation.

4.1. Study Design and Participants

The study described in this paper included five participants, three male, two female, mean age 30 ± 3 , who participated in the study for an average of 14 days. This study was carried out in a corporate research environment in which the participants each had an individual office space which was isolated from the common environment through a door. Study participants were briefed on the experiment design and participated in the experiments only after giving informed consent.

To create a uniformly sampled and labeled dataset, the user's environment was varied by changing the thermostat to a lower or higher setting, by adding a space heater, and by adding humidifier to the environment. For each parameter combination (*e.g., added heat and humidity*) the settings remained fixed for batches of four hours. However, subsequent four hour batches were alternated between hot and cold temperatures. Each user also experienced several days of environmental stability where the office HVAC controlled the room parameters.

The participants were instructed to wear the Microsoft Band¹ during work hours (9 AM to 5 PM) and to take them off only when they leave their offices for extended periods of time (*e.g., lunch, hour-long meetings; short bathroom breaks were excluded*). Participants were instructed not to manually modulate their comfort level by wearing additional layers of clothes when cold or removing layers of clothing when hot.

The features measured by the Microsoft band include heart rate, skin temperature, calories consumed, metabolic rate, altimeter, barometer, steps taken, and elevation. Data provided by the Microsoft band is labeled with user provided thermal-comfort ratings. The users were instructed to provide such ratings approximately every 20 minutes. The feedback was

provided via voice recognition with seven possible comfort levels: very cold (-3), cold (-2), chilly (-1), comfortable (0), warm (1), hot (2), and very hot (3). All data were collected via a custom Windows phone application which was written in-house to provide streaming access to the band data. Each sensor provided a sample approximately every 8 seconds.

In addition to the wearable features, room sensors were also deployed. Room temperature was recorded from three separate room thermometers (placed at different locations in the office). A gradient of the room temperature was created by subtracting instantaneous values of the temperature sensors from one another. This gradient serves as a proxy for the variation in temperature across the room and the potential motion of air mass in the room. Two of the temperature sensors, DHT11 and the NEST, also record humidity. In addition to temperature and humidity sensors, a hot-wire anemometer was used detect airflow in the room. Specifically, the hot-wire anemometer was placed near the HVAC output vent to detect room air flow. All room sensors were sampled using an Arduino connected to a Raspberry Pi. Sensor samples were collected approximately every 2-15 seconds. All feature data (room data and wearable data) was aggregated in the Microsoft Azure cloud.

4.2. Sensor Features

In the preceding section it was noted that sensors in the data collection system were sampled at different rates. The primary concern during data collection was to collect a complete training dataset. For this reason each sensor was sampled as close as possible to their maximal sampling rate as possible. The variance in sampling rates originates in technical limitations of the Microsoft Windows phone and the Raspberry Pi sensors. To capture the dynamics surrounding user provided thermal comfort ratings we extract features from time windows immediately preceding user provided ratings. Specifically we aggregate features in five minute windows preceding each thermal comfort rating. The length of the window is a parameter, but the choice of five minutes is motivated by prior work in thermoregulatory behavior (Schlader, Stannard, & Mündel, 2010).

We only extract features associated with user provided ratings that are spaced at least five minutes apart to avoid redundancy in ratings as well as to minimize noise in user supplied ratings. Table 1 provided details of the dataset including available and usable number of user provided ratings. Here the usable number of ratings represents the number of feedback reports given to the user that are spaced at least five minutes apart. In addition to the averaged features, within each five minute window we also compute the mean, variance, median, min and max for eight features: heart rate, skin temperature, core body temperature, preferred temperature, wind speed, room temperature 1, room temperature 2 and room humidity. We estimate the user's core body temperature from the user's heart

¹All users wore Microsoft band II except for one participant who wore Microsoft band I.

rate using a Kalman filter. This filter was shown in prior work to accurately track the core body temperature (Buller et al., 2013). We also derive a preferred temperature feature which uses skin and core body temperatures as follows,

$$\begin{aligned} preferredTemp = & -0.3 \times skinTemp \\ & \times (coreBodyTemp - 36.3) + 44 \end{aligned} \quad (8)$$

These features are motivated by prior work in thermal comfort estimation and prediction (Laftchiev & Nikovski, 2016). Note here that it is always possible to add more features to this type of model. The current set of features represent the features with the most significant contribution in terms of minimizing the model output error.

The statistics which are computed over five minute windows are supplemented with simple moving averages to capture trends in sensor data over short time scales. We compute simple moving averages between two to nine samples immediately preceding the user provided ratings within each five minute window. This brings the total number of features to 104.

4.3. Data Partitioning and Preprocessing

Having collected the data, an important question is how to best split the complete dataset into training and testing datasets. The optimal choice of this split is a study parameter that needs to be empirically evaluated, however for this work the labeled dataset was split into two halves for each day of the experiment and for each user. The first half is used to train and the second half is used to test the comfort prediction model.

Each collected feature is centered by subtracting the mean and dividing by the standard deviation to bring all features to the same scale. This ensures that no single feature will dominate the regression model. Both train and test datasets were transformed using the mean and standard deviation computed only on the training partition of the dataset within each user. Labeled examples were also centered using, again, normalization coefficients derived from the training data. Here only the mean was subtracted from each rating. Normalizing the labeled examples obviates the need to fit an intercept in regression settings (Hastie et al., 2009).

4.4. Active Learning - Querying Strategies

For this paper we experimented with two strategies. Each strategy is set in the pool-based active learning setting, but has been optimized for the streaming setting which is the natural extension of this work.

The first active learning strategy leverages a K nearest neighbors approach (QBC-K). The main idea of this labeling strategy is to compute the disagreement score for all available examples in the pool, \mathcal{A} . Then from this set of disagreement scores, the example chosen is that which had the maximum disagreement score. The label for this example is queried.

We compute the disagreement score as in equation (7), the first term, \hat{y}_C^i , we set C to equal K nearest neighbors. Then compute the mean rating over the K nearest neighbors, where neighbors correspond to labeled examples from $N - 1$ users and the notion of nearest is defined by Euclidean distance. The number of neighbors used in the estimate of the mean labeled examples was empirically tested for neighbor values $K = 5, 10, 15, 20$. Of these, it was observed that 10 neighbors yielded optimum performance. The second term in equation (7), $\hat{y}_{\mathcal{L}}^i$, is computed using the N^{th} user's current prediction model which is trained only using labeled examples \mathcal{L} . Specifically at budget, B , \mathcal{L} would hold at most B labeled examples, all from the N^{th} user. This strategy is a model-based querying strategy which utilizes the model of the N^{th} user. Therefore, the prediction model is retrained after each labeling point is added to \mathcal{L} .

In the second active learning strategy, each of the $N - 1$ users is treated as a committee member who is allowed to make a prediction for all available examples in \mathcal{A} . That is, for each committee member a thermal comfort model is learned using only data from that user. A 5-fold cross-validation over each user's data is performed to choose hyperparameters. Each committee member then predicts a thermal comfort rating for all available examples in the pool. Then a weighted mean of the committee ratings is computed for each sample. Higher weights are assigned to users that overlap with the N^{th} user in feature space. These weights are computed as inverse of AUROC between N^{th} user and $N - 1$ users in pairs. The remaining details of the strategy are the same as above in the first strategy. In the results section, the performance of both strategies is compared to random sampling of the available data samples.

4.5. Evaluating the Performance of the Approach

To evaluate the effectiveness of the proposed transfer active learning framework the model that is used is fixed and several decisions are made about the error reporting metrics and the format of error. This section details those decisions. The section begins with a discussion of the model choice and implementation. This is followed by a description of the error and how it is calculated. The section then concludes with four error metrics calculated throughout the framework to show the effectiveness of the proposed approach.

4.5.1. Model Selection: Ridge Regression

For this work, the model chosen was ridge regression. Note here that this model can be readily kernelized to learn a non-linear model of the data. This may be desirable in future work that focuses exclusively on determining the lowest possible error rate of the thermal comfort prediction model. In this work the focus is on demonstrating that the transfer active learning framework is capable of reducing the number of required labeled examples to achieve an error rate comparable

User	Days	No. user ratings	No. usable labels	No. non-overlap. labels (5m)	Mean duration b/w labels (5m)	Label distribution						
						-3	-2	-1	0	1	2	3
1	7	149	97	82	~27.47m	0	2	16	40	9	11	4
2	12	320	263	185	~21.72m	8	13	23	94	22	16	9
3	10	201	138	133	~30.68m	16	14	23	54	18	7	1
4	13	143	119	113	~44.69m	4	18	6	52	16	17	0
5	29	514	400	300	~48.69m	1	10	75	117	67	28	2
Total	71	1337	1017	813	–	29	57	143	357	132	79	16

Table 1. Description of the dataset including number of days and the distribution of user supplied thermal comfort ratings that are spaced at least five minutes (5m) apart

to the strictly supervised model learning approach. For clarity of presentation, the model used is linear.

In order to implement the framework a custom implementation of ridge regression was used. This in-house implementation was used to modify the OLS objective function in accordance with equation (5) to perform transfer learning. The implementation is written in Python using the `scipy` function `minimize` (Foundation, 2010) and the gradients with respect to W are manually derived and supplied within the code. Because both the design matrix X and the target values y are centered, the fitting of ridge regression is performed without fitting an intercept. The method chosen for the optimization solver is the Newton conjugate gradient method with default settings.

4.5.2. Reporting Performance

In this paper, error is reported as the mean root mean square error (RMSE) over five users. That is, an RMSE is determined for each user and the mean of the five RMSEs is reported. Each error is accompanied by a set of standard error bars. To minimize the effects of randomly seeding the analysis (*e.g.*, *cross-validation folds used to pick hyperparameters, randomly initializing weights*) each experiment is run for ten times per user. This means that each user RMSE is averaged over the ten runs before being reported.

4.5.3. Within-users (W):

A key metric to evaluate the active transfer learning framework is within-user performance. This error represents the generalization error when a model is learned (and tested) using only data from each specific user. To find the within-user error, the ridge regression model is trained and tested on training and testing data partitions which are determined as described above. For each user, a 5-fold cross-validation is performed on the training dataset to choose the optimal ridge penalty parameter, λ . The parameter range is set from $1e^{-4}$ to $1e^4$. Our observation is that the optimal parameter at no time falls on the boundary of this range.

4.5.4. Between-users (B):

To determine the starting point before active learning is employed, the between-users error shows the generalization error

of the ridge regression model when it has been trained on data from $N - 1$ users. When training this model the complete dataset for the $N - 1$ users is employed. This means that both the training and testing partitions of those users' data is used during training. Testing the model occurs on the N^{th} user's test data only. The data is preprocessed as discussed above, and a $N - 1$ -fold cross-validation is performed on the training dataset to choose the ridge penalty parameter. Here the hyperparameters follow the same range as within-user evaluation protocol.

4.5.5. Active Learning (AL) Error Trajectory:

To evaluate the performance of active learning strategies, this paper presents the trajectory of the generalization error of the ridge regression model as a function of number of training examples. To isolate the active learning trajectory from prior knowledge, the model is first initialized with a random weight vector with small magnitudes ($\pm 1e - 4$). The ridge penalty parameter is also set to a default value of $1e - 4$.

Then using equation (5) with the population level model, W_p , set to a vector of zeros to eliminate any effect of transfer learning. The reported error for each set of labeled examples, is the generalization error learned from this objective function. Active learning is performed for each user by iterating this objective after each label is received up to a given query budget B . For each user the training dataset is used as a pool of available examples and the model is tested on a held out test partitions.

As previously noted, the two query strategies employed are K nearest neighbors and the committee of $N - 1$ users. In both strategies, we pick examples with the maximum disagreement score. We compare this against a query strategy of randomly issuing queries. For these experiments, the total permissible label budget is $B = 40$ because this is the maximum number of available examples (for the user with the smallest pool of labeled examples) in the training partition given a 50%/50% train/test split. A 5-fold cross-validation is performed for each of the following label budgets, 10, 15, 20, 25, 30, 35, and 40. This cross-validation determines the optimal ridge penalty parameter using only the labeled examples for each user.

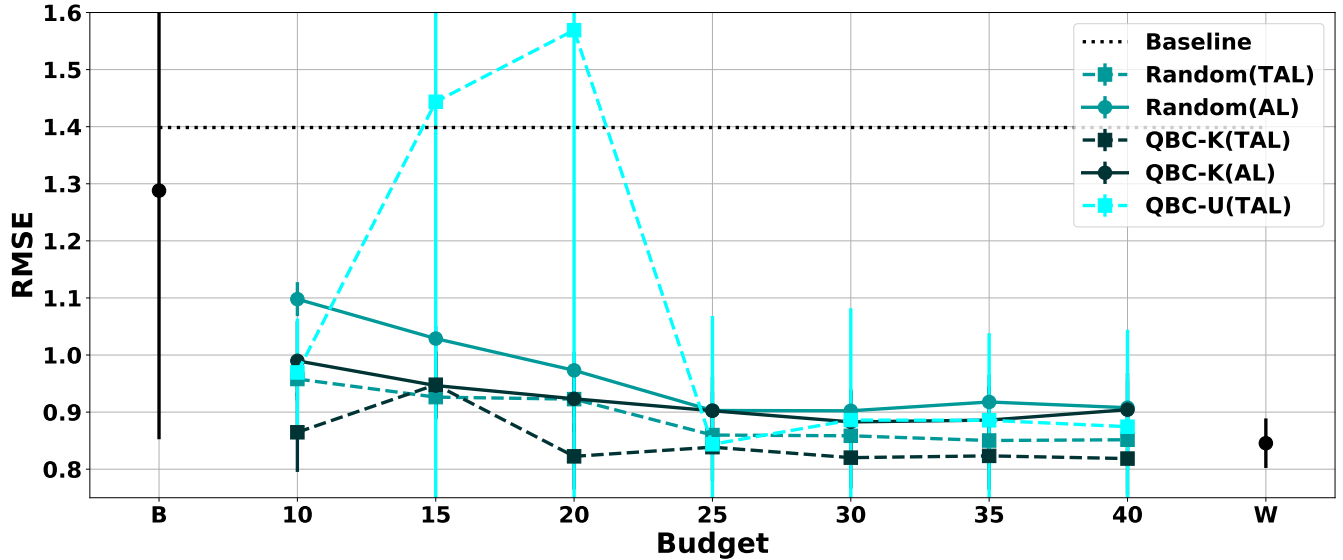


Figure 2. The performance of the Active Transfer Learning Framework in personalized thermal comfort prediction. Comparing within-user, between-user, active learning and transfer active learning evaluation protocols.

4.5.6. Transfer Active Learning (TAL) Error Trajectory:

Lastly, to evaluate the effectiveness of the transfer active learning framework, we show the trajectory of the generalization error when the between-users model is further tuned using labeled data that is actively queried from the N^{th} user. Here we use between-user model to learn the weights which is set to W_p in the objective function found in equation (5). The first query is then issued using this model. Our querying strategy is same as the one outlined for active learning above. Then for each subsequently acquired labeled example, the model is re-trained using equation (5). All other details are same as the active learning error trajectory.

5. RESULTS AND DISCUSSION

The results of the active transfer learning framework using ridge regression and the modified QBC query strategies are shown in Figure 2. This figure serves to summarize the results and to display the effectiveness of this approach. Note that all error metrics described in the previous section are included in the figure.

Figure 2 is constructed by placing the between-users (B) error rate at the extreme left of the x-axis. Here, no labeled examples have been queried from the N^{th} user. On the right extreme of the x-axis, we place the within-users (W) error rate. This error rate represents the generalization error of the model if it were to only be trained with data from the N^{th} user. In between these two evaluation protocols are plotted the results from active learning and transfer active learning. On the y-axis we show the RMSE for each of the corresponding methods. At the top of the graph, with a dotted line, we show a baseline

method which predicts that the user is always comfortable.

The figure gives the reader two important conclusions. First, transfer active learning is always better than active learning. Recall here that in active learning the model is initialized with a random weight vector. Second, the QBC-K strategy proposed in this paper outperforms other querying strategies. Notably, QBC-K outperforms the random query strategy. This means that QBC-K meaningfully adds selectivity with respect to data points which are important.

These results verify the underlying assumption of this work that there is an overlap in the thermo-regulatory behavior of users and leveraging this information leads to better performance. Furthermore, the out-performance of the QBC-K strategy as compared to random sampling, confirms the second hypothesis of this work which is that neighboring labeled examples from other users carry relevant information that can guide the personalized thermal comfort prediction model to pick and choose examples with high utility. The better performance of QBC-K can be attributed to its choice of committee members which are chosen from one or more users spanning multiple days whereas for QBC-U each of the $N - 1$ users are committee members and not all of them may be useful in computing disagreement scores.

With regard to model performance, we observe two favorable outcomes. First, the within and between subject model error evaluation shows the expected trend of a mean RMSE at 0.845 ± 0.04 and 1.288 ± 0.44 respectively. This confirms that the between subject model is more coarse than the within subject model with room for improvement. Second, all prediction models perform significantly better than the baseline model

which assumes the users are always at thermal comfort with a zero rating.

The trajectory observed for transfer active learning algorithms is mostly as expected, lower than between-subjects RMSE, with some exceptions which we will explain below. This follows intuitively from the framework development because the framework penalizes the deviation of weights from the between subject's weights. The intermittent peak in the transfer active learning trajectory is a result of hyperparameter learning that is performed in cross-validation when 15 and then 20 labeled examples are acquired respectively. The reason for this intermittent hyperparameter tuning is that the tuning proved to be very computationally expensive while RMSE suffered if tuning was not performed.

Acceptable RMSE values for untuned hyperparameters did not occur until the label budgets exceeded 50 labeled examples. For extremely low label budgets, for example < 10 , we refer to these as burn-in periods where transfer active learning used a hyperparameter tuned on $N - 1$ users and active learning used a default penalty parameter of $1e - 4$. A significant result of the transfer active learning with the QBC-K query strategy is that the resulting model is able to achieve a performance (RMSE) that is very close (0.35% difference) to the RMSE observed in the classically trained within-subject model with a small budget of 10 to 25 labeled examples. On average, these results point to a 70% reduction in labeling effort over five users. In comparison, the within-user performance was achieved with an average of 82.6 labeled examples. To further illustrate the results, the mean error across users with models trained on their respective data using the fully supervised approach (within-user) is 0.845 ± 0.04 . The transfer active learning approach proposed in this work achieves a mean error of 0.818 ± 0.05 . All errors are noted on the ASHRAE comfort scale. Here it should be noted that the model error is due to several factors including uncertainty in sensor measurements and incomplete sampling of our data space. An important topic to study in a future study is the contribution of uncertainty in sensor measurements. This is because understanding this contribution can lead us to practical suggestions about the type and grade of sensor that should be used in real systems performing online thermal comfort estimation.

We hypothesize that the improvement in performance, with respect to the number of labeled examples needed to achieve a certain RMSE value, comes from the observation that a single model trained over all data from a single user (within-user) would necessarily have a larger modeling error, especially when including any outlier labeled examples provided by the user, than a model trained on a small but relevant set of examples. However, until the development of this transfer active learning framework it was not possible to determine which examples are most important to a user on any given day. This is because one user does not provide a sufficient quantity of

user labeled examples under all possible conditions within a single day to facilitate exhaustive supervised model learning. The advent of this framework and specifically utilizing the warm start approach described in this paper, it is now possible to choose the most important data points to label within a day thereby improving the performance of the within-user model.

6. CONCLUSION

This paper demonstrated a new approach to learning personalized thermal comfort models for new user in a personalized thermal comfort prediction system. This new approach was presented as a framework that combined the machine learning fields of active and transfer learning to reduce the labeling effort needed to obtain an accurate model of thermal comfort. Importantly the approach shown here was able to reduce the labeled examples needed from new users by 70% as compared to a fully supervised approach. Specifically, the framework achieves a mean error of 0.82 ± 0.05 while the fully supervised learning approach achieves a mean error of 0.85 ± 0.04 . These results indicate a significant improvement in thermal comfort prediction at reduced quantities of user supplied labels.

REFERENCES

- Belluck, P. (2015). Chilly at work? office formula was devised for men. *New York Times*.
- Bonilla, E. V., Chai, K. M., & Williams, C. (2008). Multi-task gaussian process prediction. In *Advances in neural information processing systems (nips)* (pp. 153–160).
- Buller, M. J., Tharion, W. J., Chevront, S. N., Montain, S., Kenefick, R., Castellani, J., ... Hoyt, R. W. (2013). Estimation of human core temperature from sequential heart rate observations. *Physiological measurement*, 34(7), 781-98.
- Burbidge, R., Rowland, J. J., & King, R. D. (2007). Active learning for regression based on query by committee. In *International conference on intelligent data engineering and automated learning* (pp. 209–218).
- Cai, W., Zhang, Y., & Zhou, J. (2013). Maximizing expected model change for active learning in regression. In *International conference on data mining (icdm)* (pp. 51–60).
- Ergonomics of the thermal environment – analytical determination and interpretation of thermal comfort using calculation of the pmv and ppd indices and local thermal comfort criteria* (Tech. Rep.). (2005). ISO 7730.
- Evgeniou, T., & Pontil, M. (2004). Regularized multi-task learning. In *Acm sigkdd international conference on knowledge discovery and data mining* (pp. 109–117).
- Fanger, P. O. (1967). Calculation of thermal comfort: Introduction of a basic comfort equation. *ASHRAE Transactions*, 73.
- Farhan, A. A., Pattipati, K., Wang, B., & Luh, P. (2015). Predicting individual thermal comfort using machine

- learning algorithms. In *International conference on automation science and engineering (case)* (pp. 708–713).
- Foundation, P. S. (2010). Python language reference, version 2.7. <https://docs.scipy.org/doc/scipy/reference/generated/scipy.optimize.minimize.html>.
- Geman, S., Bienenstock, E., & Doursat, R. (2008). Neural networks and the bias/variance dilemma. *Neural Networks*, 4(1).
- Haldi, F. (2010). *Towards a unified model of occupants' behaviour and comfort for building energy simulation* (Unpublished doctoral dissertation). EPFL.
- Hastie, T., Tibshirani, R., & Friedman, J. (2009). *The elements of statistical learning: Data mining, inference, and prediction*. Springer.
- Hedge, A., Wafa, S., & Anshu, A. (2005). Thermal effects on office productivity. In *Proceedings of the human factors and ergonomics society annual meeting* (Vol. 49, pp. 823–827).
- Huang, C.-C. J., Yang, R., & Newman, M. W. (2015). The potential and challenges of inferring thermal comfort at home using commodity sensors. In *Proceedings of the international joint conference on pervasive and ubiquitous computing (ubicomp)* (pp. 1089–1100).
- IJzerman, H., & Semin, G. R. (2009). The thermometer of social relations mapping social proximity on temperature. *Psychological Science*, 20(10), 1214–1220.
- Indoor environmental input parameters for design and assessment of energy performance of buildings - addressing indoor air quality, thermal environment, lighting and acoustics* (Tech. Rep.). (2006). European Standards Commission.
- Jiang, L., & Yao, R. (2016). Modelling personal thermal sensations using c-support vector classification (c-svc) algorithm. *Building and Environment*, 99, 98–106.
- Laftchiev, E., & Nikovski, D. (2016). An iot system to estimate personal thermal comfort. In *World forum on internet of things (wf-iot)* (p. 672–677).
- Lewis, D. D., & Catlett, J. (1994). Heterogeneous uncertainty sampling for supervised learning. In *Machine learning proceedings* (p. 148 - 156). San Francisco, CA.
- Natarajan, A., Gaiser, E., Angarita, G., Malison, R., Ganesan, D., & Marlin, B. (2014). Conditional random fields for morphological analysis of wireless ecg signals. In *Acm conference on bioinformatics, computational biology, and health informatics* (pp. 370–379).
- O'Neill, J., Delany, S. J., & MacNamee, B. (2017). Model-free and model-based active learning for regression. In *Advances in computational intelligence systems* (pp. 375–386). Springer.
- Pan, S. J., & Yang, Q. (2010). A survey on transfer learning. *IEEE Transactions on Knowledge and Data Engineering*, 22(10), 1345–1359.
- Ranjan, J., & Scott, J. (2016). Thermalsense: Determining dynamic thermal comfort preferences using thermographic imaging. In *Proceedings of the international joint conference on pervasive and ubiquitous computing (ubicomp)* (pp. 1212–1222).
- Schlader, Z. J., Stannard, S. R., & Mündel, T. (2010). Human thermoregulatory behavior during rest and exercise - a prospective review. *Physiology & behavior*, 99(3), 269–275.
- Schwaighofer, A., Tresp, V., & Yu, K. (2005). Learning gaussian process kernels via hierarchical bayes. In *Advances in neural information processing systems (nips)* (pp. 1209–1216).
- Settles, B. (2010). *Active learning literature survey* (Tech. Rep.).
- Sugiyama, M. (2006). Active learning in approximately linear regression based on conditional expectation of generalization error. *Journal of Machine Learning Research*, 7, 141–166.
- Thermal environmental conditions for human occupancy* (Tech. Rep.). (2013). ASHRAE.
- Wang, X., Huang, T.-K., & Schneider, J. (2014). Active transfer learning under model shift. In *International conference on machine learning (icml)* (pp. 1305–1313).

BIOGRAPHIES



Annamalai Natarajan Annamalai Natarajan is a research scientist with Philips Research. His research interests are in machine learning for healthcare with emphasis on physiological time series. He holds a PhD in computer science from UMass Amherst, a M.S. in computer science from Colorado State University, Fort Collins and a B. Tech. in Information Technology from University of Madras, India.



Emil Laftchiev Emil I. Laftchiev received the B.S. and Ph.D. degrees in Electrical Engineering from the Pennsylvania State University in 2007 and 2015, respectively. In 2015 Emil joined MERL in Cambridge, MA, USA as a research scientist. He has a wide range of research interests including machine learning, the Internet of Things, time series analysis, and signal processing. Within these fields Emil is particularly interested in developing subfields of machine learning such as transfer learning, active learning, and generative modeling. Emil is a member of IEEE and SIGKDD. His work has appeared in IEEE Conferences, Journals and Workshops.