Building Technologies & Urban Systems Division
Energy Technologies Area
Lawrence Berkeley National Laboratory

Sharing is caring: An extensive analysis of parameter-based transfer learning for the prediction of building thermal dynamics

Giuseppe Pinto[1], Riccardo Messina[1], Han Li[2], Tianzhen Hong[2], Marco Savino Piscitelli[1], Alfonso Capozzoli[1]

[1]Department of Energy, TEBE Research Group, BAEDA Lab
[2]Building Technology and Urban Systems Division, Lawrence Berkeley National Laboratory

Energy Technologies Area
October 2022

DOI: 10.1016/j.enbuild.2022.112530

This work was supported by the Assistant Secretary for Energy Efficiency and Renewable Energy, Building Technologies Office, of the US Department of Energy under Contract No. DE-AC02-05CH11231.

Sharing Is Caring: An extensive analysis of parameter-based transfer learning for the prediction of building thermal dynamics

Giuseppe Pinto[a,b], Riccardo Messina[a], Han Li[b], Tianzhen Hong[b], Marco Savino Piscitelli[a], Alfonso Capozzoli[a]

[a]*Department of Energy, TEBE Research Group, BAEDA lab, Politecnico di Torino, Italy*
[b]*Building Technology and Urban Systems Division, Lawrence Berkeley National Laboratory, One Cyclotron Road, Berkeley, CA 94720, United States*

Abstract

In recent years deep neural networks have been proposed as a lightweight data-driven model to capture high-dimensional, nonlinear physical processes to predict building thermal responses. However, the need of a large amount of data for the training process of deep neural networks clashes with the potential limited data availability in most existing or new buildings. Transfer learning aims to improve the performance of a target learner exploiting knowledge from related and similar environments. This study conducted a suite of experiments that leveraged 250 data-driven models based on a synthetic dataset of a building archetype to study the influence of data availability, energy efficiency level, occupancy and climate for the transfer process of thermal dynamics. The performance of the transfer learning process was compared against a classical machine learning approach. The results suggest that building thermal dynamics can be effectively transferred under the same climatic conditions, increasing performance when dealing with different occupancy schedules, efficiency levels and low data availability. Furthermore, the paper compares the performance of both transfer learning and machine learning approaches in an online fashion, to support the implementation in real-world deployment.

Keywords: Transfer Learning, Building Thermal Dynamics, Deep Neural Network, Data-driven Models, Grid-interactive Buildings

1. Introduction

The current energy transition is deeply changing the way energy is used and generated. The need of a further decarbonisation of the building sector [1] has fostered the use of distributed renewable energy resources. In this framework, Grid-interactive Efficient Buildings (GEB) [2] play a key role in the energy transition, exploiting advanced control strategies, identified as a way to increase energy savings up to 28% [3] while providing benefits for

*Corresponding author
Email address: alfonso.capozzoli@polito.it (Alfonso Capozzoli)

the electric grid [4]. However, the main bottleneck for their widespread implementation is that these control strategies often rely on predictive-based optimization, which requires the development of a fast and accurate building thermal dynamic model [5]. A thermal dynamic model of a building (usually built at the resolution of a thermal zone or room/space) predicts how the indoor environmental conditions (e.g., indoor air temperature and humidity) respond to the internal disturbances (e.g., heat gains from occupants, lighting and plug-in equipment), external factors (e.g., outdoor air temperature, humidity, solar irradiance), and HVAC operations (zone temperature setpoint, supplied cooling or heating energy). The thermal dynamics of a building is governed by high-dimensional, nonlinear and discontinuous dynamics, which require effort and expertise to be proper modeled [6].

In particular, three main techniques are used to model building thermal dynamics: white-box modeling, gray-box modeling and black-box modeling. White-box models use physical knowledge to describe building dynamics [5] and are based on the concept of heat transfer and energy and mass conservation. The major barrier of white-box modeling is represented by the time and effort necessary to define and collect reliable building features. The gray-box category covers a wide range of models that exploit simplified physical relationships but also require parameter estimation based on measured data. A typical concept in gray-box modeling is the RC analogy that defines any model by its affinity with a resistor-capacitor electrical circuit [7], and their challenge is related to a robust parameter estimation method. Black-box models learn the building thermal dynamics directly from the measured historical data, without making prior assumptions regarding any physical relationships [8]. The main advantages of the black-box models are the lower development cost and the flexibility in using any measured signal as an input or output, due to the absence of physics involved. On the other hand, black-box models require a large amount of training data and are not reliable outside the training range. Due to the higher complexity of white-box models and the increasing adoption of sensing and metering in buildings that provide a growing amount of building-related data, there is a rising trend towards the use of data-driven methods for advanced control in buildings [9, 10].

1.1. Literature review

The present subsection introduces a literature review on the application of black-box models for the building thermal dynamics prediction, highlighting common practices and limitations while identifying innovative approaches in the field. Among the first applications, Ruano et al. [11] explored the use of a radial basis function neural network to predict the indoor air temperature of a public building. Sun et al. [12] proposed a multiple linear regression model to predict the supply temperature of a district heating network, adjusting it according to actual indoor temperature deviation. Shi et al. [13] used a back-propagation neural network to predict indoor relative humidity and air temperature with different forecasting horizon. Kusiak and Xu [14] proposed a dynamic neural network to relate HVAC energy consumption with indoor temperature evolution, optimizing the control strategy of the HVAC system with a data-driven approach. Similarly, nonlinear [15, 16] autoregressive neural networks for the indoor temperature prediction were integrated with controllers to optimize the HVAC systems. More recently, Huang et al. [17] implemented a predictive

controller coupled with a neural network able to estimate the indoor air temperature of a multi-zone building, to optimize the start and stop of an HVAC system, while Drgoňa et al. [10] exploited neural networks and regression trees to construct an approximate model predictive controller.

To evaluate the effectiveness of the different machine learning (ML) techniques and neural network architecture, Wang et al. [18] applied 12 data-driven models with the aim of predicting the heat load of a single building. Results showed Long Short Term Memory (LSTM) and eXtreme Gradient Boosting (XGBoost) to be the best, respectively, for short-term load prediction and long-term load prediction.

Recently, much interest has been devoted to the application of LSTM for short-term load and thermal dynamic prediction. In [19] an LSTM neural-network was employed to predict the indoor air temperature in a multi-zone building. Xu et al. [20] compared two LSTM models to predict indoor temperature evolution one step ahead and multiple time steps ahead, studying the advantages of using an error correction for multiple time steps ahead. Ellis and Chinde [21] used an Encoder-Decoder LSTM to describe the dynamic of an air-handling unit with variable air volume relating it to the indoor temperature evolution, coupling the information with a model predictive controller (MPC) to reduce energy costs. Fang et al. [22] proposed three LSTM-based sequences to sequence model architectures to perform multi-step ahead indoor air temperature forecasting: an LSTM-Dense model, an LSTM-LSTM model and an LSTM-dense-LSTM model, evaluating the performance under different forecast horizons. The results and analyses showed that the LSTM-dense model performs better for shorter forecast horizons, while the other two are more suitable for longer forecast horizons. Lastly, Pinto et al. [23] proposed LSTM models to predict indoor temperature evolution in multiple buildings, deploying one model per building and coupling them with a deep reinforcement learning (DRL) controller to perform district demand side management, speeding thermal dynamics simulation at the district level. Recently, a new paradigm in neural network was introduced with physics-informed neural networks (PINNs) [24]. These neural networks are trained to solve supervised learning tasks while respecting any given laws of physics described by general nonlinear partial differential equations, combining the advantages of white-box modeling with black-box modeling. However, despite the interest in this field, only a few works explored PINNs in the domain of building energy control. Bünning et al. [25] compared physics-informed Autoregressive-Moving-Average with Exogenous Inputs (ARMAX) models to Machine Learning models based on Random Forests and Input Convex Neural Networks. In [26] a physics-informed neural network was used to predict temperature evolution in a building, increasing the sample efficiency of neural networks and performances for longer prediction horizons. The authors of [27] introduced a physics-constrained recurrent neural network (RNN) to model the thermal dynamics of buildings constraining the eigenvalues of the model, and using penalty methods to impose physically meaningful boundary conditions to the learned dynamics. Di Natale et al. [28] proposed a physics-informed NN that predicts indoor air temperature with respect to different control inputs, zone-zone, and outdoor-zone air temperature differences. However, although Gokhale et al. [26] proved a greater sample efficiency of PINNs, they still need a lot of data and physics knowledge. Furthermore, collecting and preparing a large amount

of high quality data to train machine learning algorithms is time consuming and not always feasible.

To address this gap, one key technique needed is the possibility to transfer machine learning models from one building to another, a concept called transfer learning (TL) [29]. Pinto et al. [30] reviewed the application of transfer learning in smart buildings, highlighting the role of transfer learning for building thermal dynamics models, the challenges and the future directions. However, the application of TL for building thermal dynamics is still in its infancy, so only a few studies have been conducted. Of those, none have thoroughly analyzed under which circumstances transfer learning performs better than classical machine learning, therefore, a thorough analysis is required. Hossain et al. [31] trained a Bayesian neural network (BNN) to directly learn an RC model rather than estimating parameters. The work proved that at least several weeks of data are necessary to obtain good performance for classical ML models. The paper proposes a methodology to overcome extreme data scarcity (one day of data) by identifying the best RC model based on historical consumption patterns that outperforms time-series methods that were directly built using available data. Jiang and Lee [32] pretrained an LSTM sequence-to-sequence model using a large amount of data from source buildings to study building temperature evolution and used this knowledge to fine-tune a model enhancing the performance of the target building, while [33] froze neural network hidden layers to transfer the building thermal model. Similarly, Chen et al. [34] applied transfer learning to predict not only internal temperature but also relative humidity. Lastly, Grubinger et al. [35] presented an interesting approach of online transfer learning, integrating the prediction with an MPC controller, paving the way for real-world applications.

1.2. Research gaps and contributions

Despite the opportunity to overcome the data availability issue presented for the classical ML problems, the literature review [30] highlighted the following research gaps:

1. It is still unclear how to identify the source building to perform effective transfer learning.
2. Further studies are necessary to quantify the most important features of transfer learning for building dynamics.
3. Minimum data availability for an effective transfer and the time-horizon applications have not been fully explored.

In this paper, a synthetic dataset was used to create multiple energy models of a single building in different conditions, changing building features such as efficiency level, occupancy and climate. The dataset was then used to train and compare machine learning and transfer learning models. A machine learning model only leverages data available for the target building, while the transfer learning model reuses knowledge from a source building to reduce implementation costs, speed up the training and increase performance. The aim is to assess their performance, isolating the contribution of specific features and studying the effect of data availability on transfer learning performance. With this in mind, this study aimed to address the literature gaps, with the following contributions:

1. Isolating and evaluating the contribution of key features in determining machine learning and transfer learning effectiveness, using a synthetic building dataset gathered from a detailed physics-based building energy model.
2. Performing a statistical investigation by developing approximately 250 models to assess the feature importance and data availability impact.
3. Conducting a specific analysis of negative transfer to assess the limitations of transfer learning for building thermal dynamics, to identify guidelines for future research.
4. Assessing the effectiveness of transfer learning in an online deployment setting, supporting its real-world implementation

2. Methods

The following sections describe the methods used within the paper, starting with the theoretical background of transfer learning, followed by a brief description of the neural networks exploited during the analysis.

2.1. Transfer learning

The starting point for the definition of transfer learning is the description of the concepts of "domain" and "task," reported below according to Pan and Yang [29].

Definition 1. Domain: a domain $\mathscr{D}$ consists of two components: a feature space $\mathscr{X}$ and a marginal probability distribution $P(X)$, where $X = \{x_1, \ldots, x_n\} \in \mathscr{X}$.

For example, if the learning task is the building thermal dynamic prediction, modeled as a regression problem, then $\mathscr{X}$ is the space of all influencing variables, (e.g., external temperature, occupancy, HVAC load), while x_i represents the i^{th} influencing variables and X a specific learning sample.

Definition 2. Task: a task consists of two components: a label space Y and an objective predictive function $f(\cdot)$ (denoted by $\mathscr{T} = \{Y, f(\cdot)\}$), which is not observed but can be learned from the training data, represented by a pair $\{x_i, y_i\}$, where $x_i \in \mathscr{X}$ and $y_i \in \mathscr{Y}$. The function $f(\cdot)$ is used to approximate the conditional probability $P(y|x)$ and predict the corresponding label of a new instance x.

Considering the same application of building thermal dynamic prediction, Y is a continuous space with the possible values of the internal (indoor air) temperature.

Lastly, the transfer learning definition is provided. To ease the comprehension, the definition only considers the case of one source domain $\mathscr{D}_S$ and one target domain $\mathscr{D}_T$, since it represents the most common research problem. In particular, we denote the source domain data as $D_S = \{(x_{S1}, y_{S1}), \ldots, (x_{Sn_S}, y_{Sn_s})\}$, where $x_{Si} \in X_S$ is the data instance and $y_{Si} \in Y_S$ is the corresponding output. Similarly, the target domain data are denoted as $D_T = \{(x_{T1}, y_{T1}), \ldots, (x_{Tn_T}, y_{Tn_T})\}$, where $x_{Ti} \in X_T$ and $y_{Ti} \in Y_T$ are the corresponding outputs. In many cases, transfer learning provides advantages where $0 \leq n_T \ll n_S$.

Definition 3. Transfer Learning: Given a source domain $\mathscr{D}_S$ and learning task $\mathscr{T}_S$, a target domain $\mathscr{D}_T$, and a learning task $\mathscr{T}_T$, transfer learning aims to help improve the learning of the target predictive function $f(\cdot)$ in $\mathscr{D}_T$ using the knowledge in $\mathscr{D}_S$ and $\mathscr{T}_S$, where $\mathscr{D}_S \neq \mathscr{D}_T$, or $\mathscr{T}_S \neq \mathscr{T}_T$.

Transfer learning can be classified according to label availability, domain and task similarity and technique used to transfer the knowledge.

Looking at label availability, there are three main categories: inductive, transductive and unsupervised transfer learning.

- In inductive transfer learning, both the source and target domains have labeled data, yet the source and target tasks are different.
- In transductive transfer learning, the source and target tasks are the same, yet the source and target domains are different. In this setting, the source domain has sufficient labeled data while the target domain has none.
- In unsupervised transfer learning the settings are similar to that in inductive learning, i.e., the source and target domains are the same with different but related tasks. However, there are no labeled data in both domains, and the aim is to explore the intrinsic data characteristics in different domains.

Moving to the domain and task similarity, there are mainly two cases: i.e., heterogeneous and homogeneous transfer learning. In the space classification, if the feature space and the label space of both source and target domain are the same, the scenario is classified as *homogeneous transfer learning.* Otherwise, if they have a different feature space or label space, the scenario is classified as *heterogeneous transfer learning.* Lastly, transfer learning can also be categorized based on the strategy adopted to share the knowledge, i.e., data instance-based, model parameter-based, feature representation-based, and relational knowledge-based strategies; the classification based on the strategy adopted hereafter will be defined as *solution classification.*

- The instance-based TL exploits data from a source domain in a target domain. The reasoning behind this is that a subset of data from the source domain can be used to improve the task in the target domain. To incorporate source domain data into the target task training process, one common practice is to use re-weighting and importance sampling techniques. These techniques are typically used when the data variables are the same across different domains, and they increase the amount of data available for training without substantially changing the algorithm itself.
- The feature representation-based TL extracts and exploits features to map instances from the source and target domains to improve training on the target task. A popular approach is to identify a latent feature space from the source domain, based on which the marginal distributions between two domains are minimized.

- Relational knowledge-based TL is generally used with multi-relational datasets. The underlying assumption is that some relationship among the data in the source and target domains are similar. Thus, the knowledge to be transferred is the relationship among the data.

- The model parameter-based TL assumes that the source tasks and the target tasks share some parameters or prior distributions of the hyper-parameters of the models (e.g., neural networks). The latter is based on the assumption that models developed for similar tasks will have similar model parameters or hyper-parameters. The knowledge gained from the source task is transferred to another task as shared model weights in this case. The recent success of deep learning has spawned a new type of transfer learning — network-based transfer learning [36] — which belongs to the parameter-based transfer learning category and can be further classified based on the strategy used to share model parameters:

  - The first way is to use the pretrained model for feature extraction. In this case, the weights of some layers are fixed, except for the input/output layer, which are domain dependent and need to be fine-tuned using target data. The main advantage is represented by the reduced amount of data needed to train the model, as well as the possibility to exploit data from different domains.
  - The second way is to use the pretrained model for weight initialization and fine-tuning. In such a case, the weights of the pretrained model are used for initialization purposes only and can be adjusted through a fine-tuning process.

Figure 1 displays the two strategies adopted to perform parameter-based TL, henceforth called *feature-extraction* and *weight-initialization.* Considering that in this work different neural networks share the same feature and label space, a homogeneous inductive problem using model parameter-based TL was explored.

2.2. Deep Neural networks

The study employed both the Multi-Layer Perceptron (MLP) and Long Short-Term Memory networks (LSTM). An MLP consists of at least three layers of nodes: an input layer, a hidden layer and an output layer. Except for the input nodes, each node is a neuron that uses a nonlinear activation function. MLP utilizes a supervised learning technique called *backpropagation* for training. However, classical MLPs are finite response systems that are unable to capture long-term temporal dependencies. On the other hand, LSTM are a special kind of recurrent neural network, capable of learning long-term dependencies [37]. This property of LSTMs is due to a particular gating mechanism and to the presence of two states:

- Hidden state: responsible for maintaining the short-term memory

- Cell state: responsible for maintaining the long-term memory

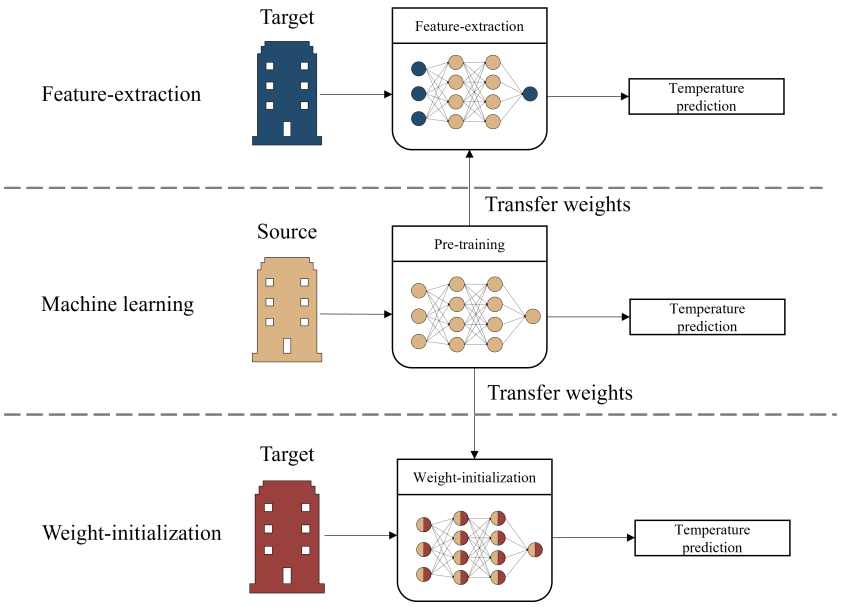

Figure 1: Parameter-based transfer learning further classified in feature-extraction (top) and weight-initialization (bottom)

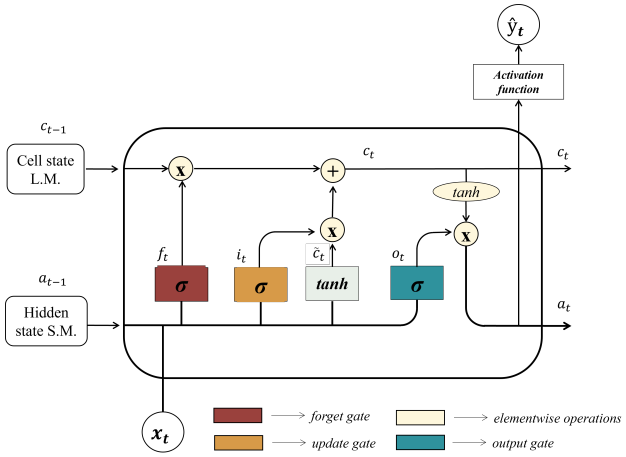

Figure 2: A schematic representation of a Long-Short Term Memory neural network

The scheme of the LSTM cell is shown in Figure 2:

The main feature of LSTMs is the cell state, which is responsible for maintaining long-term dependencies: information is removed or added to the cell state by means of three gates. The *forget* gate decides what information has to be deleted from the cell state, the *update* gate decides which information is going to be stored in the cell state, and the *output* gate is used to compute the output of the LSTM cell.

3. Case study

The selected case study is an archetype building energy model developed from the U.S. Department of Energy (DOE). The model is a medium-sized office building with three floors and a total floor area of 4,890 square meters [38]. The building consists of 12 space types: open and enclosed office rooms, conference room, classroom, dining area, lobby, corridor, stair, storage, restroom, plenum and mechanical room. A schematic representation of a floor is shown in Figure 3.

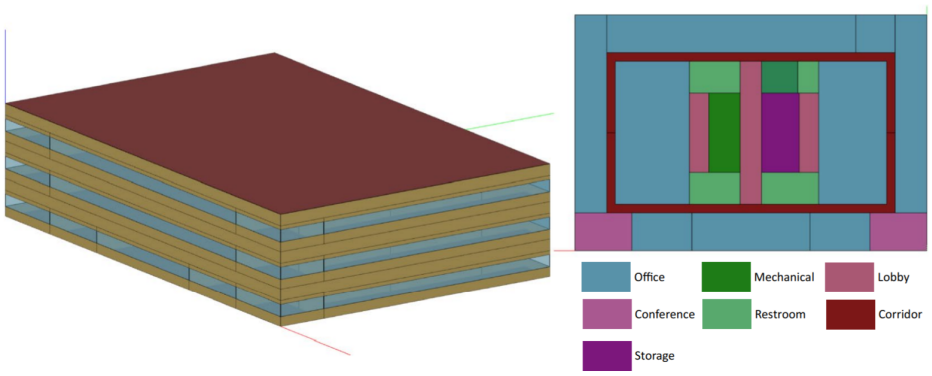

Figure 3: A schematic representation of medium office geometry and thermal zones for a single floor

The synthetic dataset includes simulations for the selected building model in different climates for multiple years, efficiency levels and occupancy patterns. The dataset includes three energy efficiency levels, obtained by changing building envelope properties, and the efficiencies of lighting, miscellaneous electric loads (MELs) , and HVAC systems. Furthermore, three sets of schedules for zone-level occupancy, lighting, MELs, and thermostat setpoint, reflecting realistic building operations from stochastic occupancy simulations, were used [39]. The resulting configuration are reported in Table 1.

Table 1: Parameters and modified features used for the design of experiment

Parameter	Cases	Features involved
Efficiency	Low, Standard, High	Building envelope properties, efficiency of lighting, MELs and HVAC systems
Climate	1A,3C,5A	Outdoor air temperature, solar radiation
Occupancy	1,2,3	Schedule of occupancy, MELs, lighting, setpoints

To study the contribution of different weather conditions on model performance, three typical climate zones were selected: Miami (1A, hot and humid), San Francisco (3C, moderate/mild), and Chicago (5A, cold winter and hot summer). A synthetic dataset [40] was used with twofold advantages: (i) it can reflect the effects of different influencing variables on building operation, and (ii) it isolates the contribution of specific features on machine learning and transfer learning model accuracy.

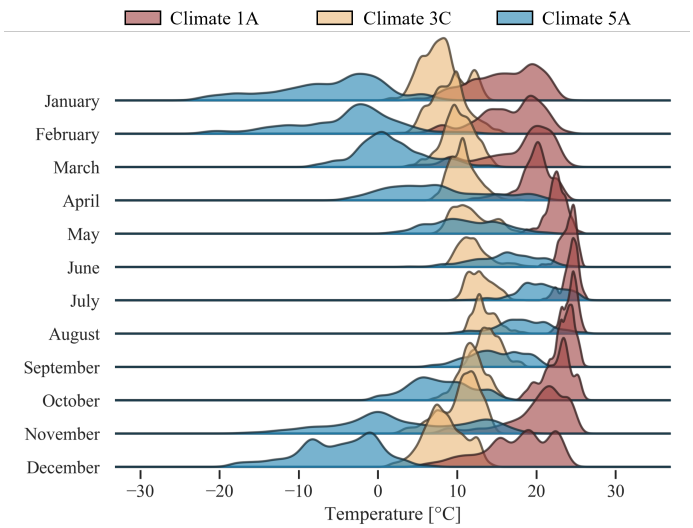

Figure 4: Distribution of the outdoor air temperature for each month and climate considered during the analysis

Figure 4 shows the outdoor air temperature distributions for the three climates selected. The selected climatic zones exhibit very different temperature patterns, with Climate 1A being cooling dominant, Climate 5A being heating dominant, and Climate 3C representing a mild climate. For each combination between efficiency level, occupancy and climate, up to two years of meteorological data were used for training and testing purposes. The simulations yielded time-series data that included whole-building and end-use energy metering, indoor and outdoor environmental parameters, and system and component variables (e.g., zone thermostat setpoints, VAV terminal supply air temperature). For a detailed description of how the synthetic dataset was obtained, refer to Li et al. [40].

4. Methodology

This section reports the methodological framework adopted, as shown in Figure 5. The methodology unfolds in four main steps, described below.

4.1. Source building selection

The first step consists in the identification of the "source building," used as a starting point for transfer learning. As pointed out in the previous section, the dataset analyzed refers to a medium-sized office building simulated in 3 climates, 3 energy efficiency levels, 3 stochastic occupancy schedules, for a total of 27 EnergyPlus models. The source building was conceived with a standard energy efficiency, the occupancy profile 1 (according to Table 1) and was simulated in Climate 3C. The climate and the energy efficiency level were chosen to represent an intermediate condition between the other two options, with the aim to further evaluate the potential of applying transfer learning. The dataset has a 10-minute granularity, with information related to whole building variables as well as zone variables.

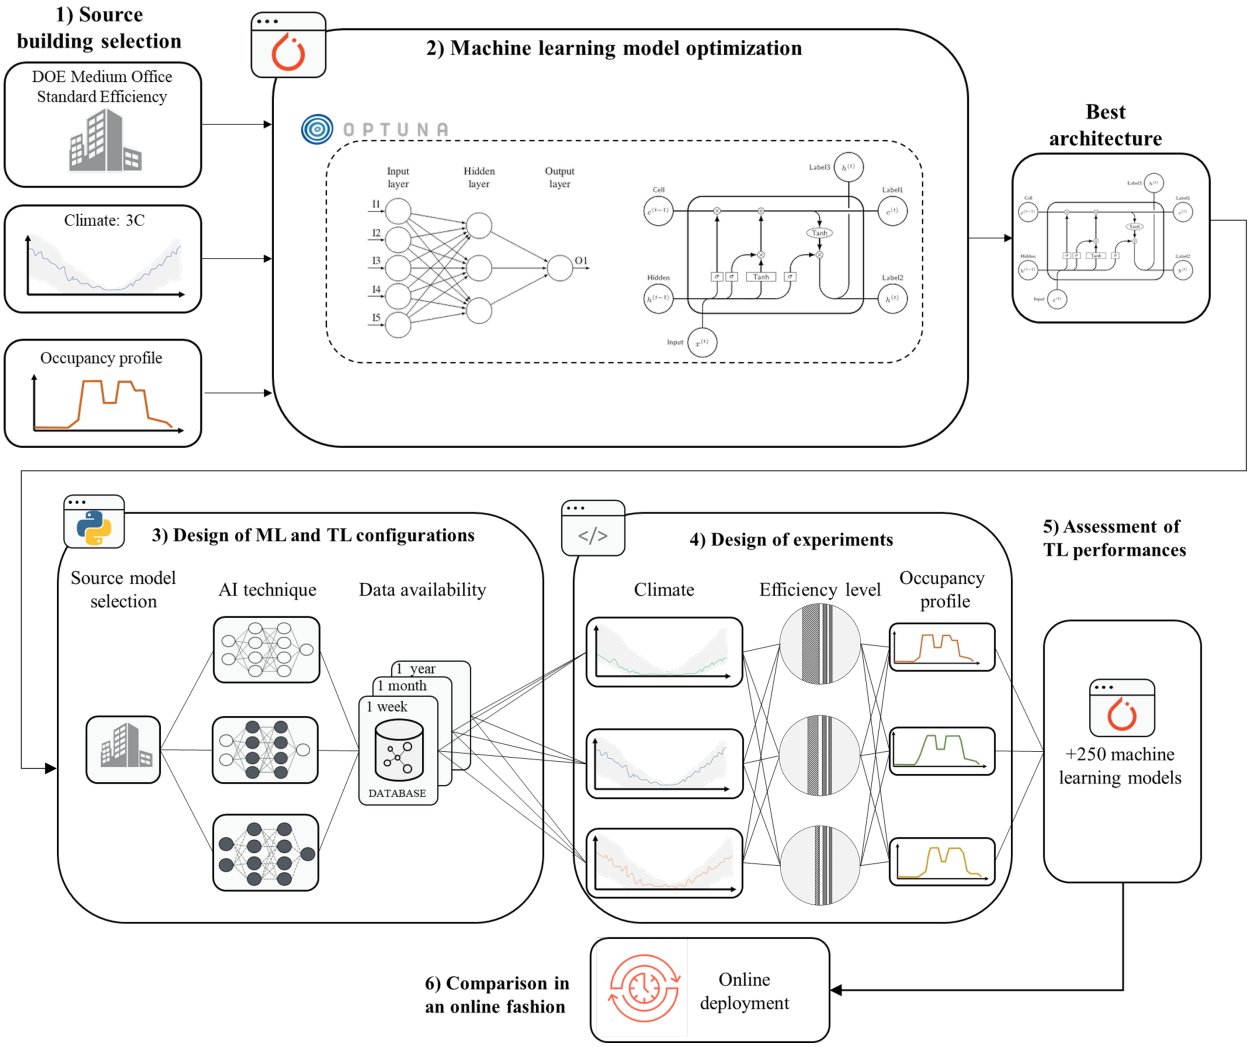

Figure 5: Methodological framework

4.2. Machine learning model optimization

The second step includes the model development, the selection of the architecture and the optimization of the related hyperparameters. The models aimed to predict the temperature evolution of a single zone (mid-office) one-hour ahead (six time-steps), exploiting information of the specific zone. This was necessary due to the impossibility of aggregating data at a higher level, since different zones may have different setpoints and occupancy schedules. The selected inputs for the machine learning models were the zone heating and cooling temperature setpoints, the outdoor air temperature, the previous internal (zone air) temperature, solar radiation, and information about hour, day and month. Figure 6 shows the input parameter together with the sliding window approach used to perform the predictions.

The architectures selected were MLP and LSTM. The developed models used 48 time-

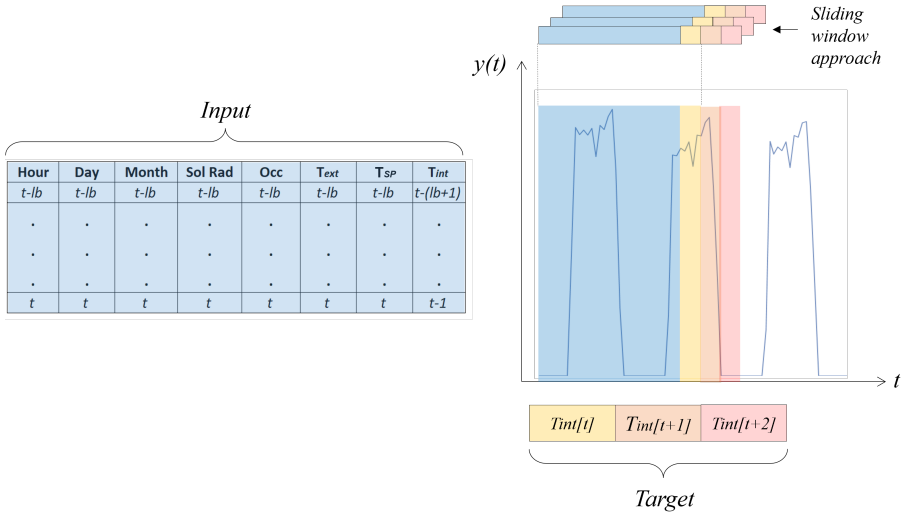

Figure 6: Input of the neural networks and sliding window approach

Table 2: Neural network hyperparameter optimization process

MLP	Range	Optimum	LSTM	Range	Optimum
# Neurons layer 1	[50-200]	100	# LSTM layer	[3-7]	3
# Neurons layer 2	[50-150]	70	# Neurons per layer	[70-300]	175
# Neurons layer 3	[20-90]	70	Epochs	[80-120]	90
# Neurons layer 4	[10-70]	10	Learning rate	[7-8.5*10^{-3}]	7.7*10^{-3}
Epochs	[80-200]	120	Batch size	[800-1000]	900
Learning rate	[7-8.5*10^{-3}]	7.57*10^{-3}	Optimizer		Adam
Batch size	[800-1000]	900			
Optimizer		Adam			
MAPE		1.096			0.535

steps (8 hours) as a lookback period to predict the next 6 time-steps (1 hour). Each architecture was characterised by specific hyperparameters, therefore an optimization process was carried out using the Optuna [41] framework. The tool allows the optimal hyperparameter combination to be searched by performing an automatic grid-search. The work performed the grid-search using five values in the specified interval shown in Table 2 with a uniform distribution. The dataset included two years of data: one used for training and validation and the other one used for testing. Table 2 illustrates the hyperparameters subject to the grid-search optimization with their optimized values, as well as the value of the mean absolute percentage error (MAPE) evaluated in the testing period. Table 2 highlights the higher accuracy of the LSTM architecture, which was then selected to perform the experiments.

4.3. Design of ML and TL configurations

The third step compares classical ML with two TL techniques to predict indoor air temperature evolution. A classical machine learning approach used the optimal hyperparameter identified in step 2 to train LSTM models on data available for the target building. The performance of the LSTM model was then compared with that resulting from the models trained using two transfer learning methods: weight-initialization and feature extraction. In weight-initialization, the whole network is fine-tuned using the data available in the target building and a lower learning rate with respect to the one used to train the source network, while in the feature extraction, the LSTM layers are frozen and only the last dense layer is fine-tuned. For both weight-initialization and feature-extraction, a learning rate equal to $2 * 10^{-3}$ was used to train the LSTM for 80 epochs.

Moreover, this step aims to analyse the impact of data availability on model performance. To this purpose, three cases were considered regarding the data availability for the target building: (i) 1 week of data, (ii) 1 month of data, and (iii) 1 year of data. The cases of one week and one month of data were used to represent a data-scarcity context and had the main purpose of highlighting in which conditions TL performs better than ML and the minimum amount of data necessary to develop an effective ML model. On the other hand, an ideal case that considered one year of data available in the target building was used to assess the generalizability of TL over ML, to assess if TL can provide additional advantages even in the presence of an extensive amount of data for the target building.

4.4. Design of experiments

The fourth step deals with the design of the scenarios resulting from the combination of the different features for the target building as reported in Section 3. Machine learning and the two transfer learning strategies were implemented to consider the combination of three climates, three energy efficiency levels, three occupancy patterns and three data availability periods. This led to 243 different models, including the one related to the source model used for transfer learning. These simulations were used to perform a statistical investigation on the most important features for the application of TL for building thermal dynamic models. All the information on the data, the code and the results produced by the statistical investigation are open-source and available at the following link: `https://github.com/baeda-polito/Transfer_learning_building_dynamics`.

4.5. Assessment of TL performance

Lastly, model performance is compared using several metrics. In particular, model absolute performance was compared using metrics such as MAE , MAPE, and MSE, the definition of which is provided below. Relative performance was quantified using the asymptotic performance and jumpstart.

$$MAE = \frac{1}{n}\sum_{i=1}^{n}|\hat{y}_i - y_i| \tag{1}$$

$$MAPE = \frac{1}{n}\sum_{i=1}^{n}\left|\frac{\hat{y}_i - y_i}{y_i}\right| \tag{2}$$

$$MSE = \frac{1}{n}\sum_{i=1}^{n}|\hat{y}_i - y_i|^2 \tag{3}$$

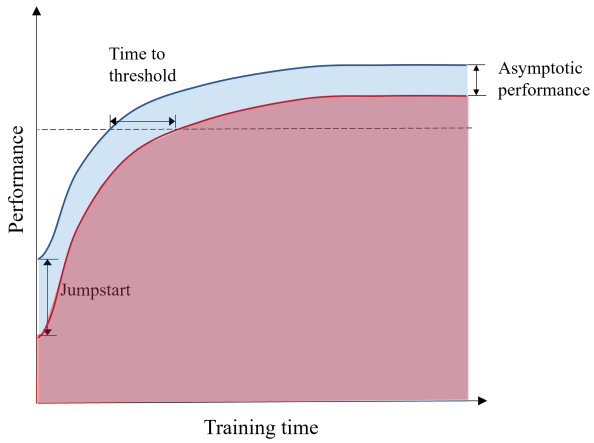

Figure 7: Common measures used to quantify the performance of transfer learning

Figure 7 displays three common measures used to assess the learning improvement. The first metric is *jumpstart,* which describes the increase in the initial performance achievable in the target task using the transferred knowledge, before any further learning. The second metric is the *time to threshold,* which is used to assess the amount of time the model takes to achieve certain performance in the target task given the transferred knowledge, compared to the amount of time necessary to learn it from scratch. The last metric is *asymptotic performance level,* which quantifies the ultimate performance level of the agent in the target task when transferred from the source. Regardless of the specific ML task, jumpstart and asymptotic performances for the regression problem are evaluated using MAE, MAPE and MSE. However, in this paper the metric time to threshold was not analysed due to the necessity to quantify a specific threshold (e.g., MAE = 0.5 °C), which may or may not ever be reached by machine learning models.

4.6. Comparison in an online fashion

To further demonstrate the effectiveness of transfer learning in an online fashion, this study compared an online machine learning approach (updating the weights of the neural network as new data become available) with an online transfer learning deployment strategy. Online transfer learning leverages one year of source data and updates the model in an online fashion each week as new data become available, performing a fine-tuning of the model. The comparison is helpful since real-world application often works with online data and building thermal dynamic models are used as a part of a model predictive control implementation, thus requiring it to be robust and fast.

5. Results

This section describes and analyses the results obtained from the proposed design of experiments. Section 5.1 describes the results obtained from both ML and TL models, analysing the performance distribution and identifying the factors that most influence model performance. Furthermore, statistical analysis was performed to compare absolute and relative performance of the proposed approaches with respect to the different features. Section 5.2 focuses on negative transfer, describing the boundary conditions in which it occurs and assessing benefits and limitations. Lastly, Section 5.3 describes computational advantages related to the application of TL, analysing jumpstart and training asymptotic performance.

5.1. Machine learning and transfer learning performance

Figure 8 shows the average performance over the entire design of experiments of ML and TL models using one month of data to assess the previously introduced metrics (MAE, MSE, MAPE) over all the six time-steps. As can be seen, the ML algorithm error is almost constant over the time-steps, while both transfer learning techniques show a lower error for the first prediction time-step, reaching about the same accuracy at the last time-step (one hour). On average, both feature extraction and weight initialization techniques perform better than machine learning. The analysis of MAE, expressed in °C shows that for the first time-step the two TL techniques have a value of 0.17 °C smaller compared to standard ML, achieving a performance improvement of 50%. Similar considerations can be made for the other two metrics, that show substantial improvement with respect to ML performance for the first time-step and a better average performance.

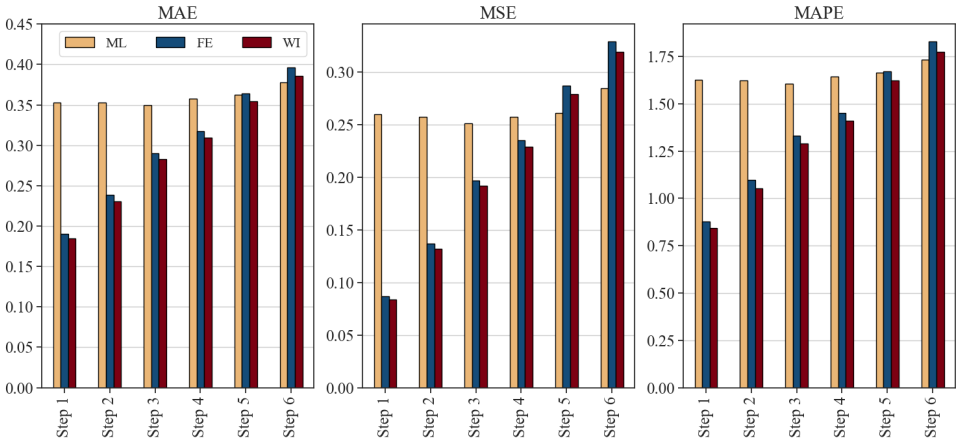

Figure 8: Performance of the different techniques over the control horizon

For the sake of simplicity, the following analysis considers only the average performance of the mean absolute error over the entire prediction horizon, since it can be interpreted easily.

The first step aimed to assess the effectiveness of transfer learning between different zones of the building. To prove the effectiveness of TL in different zones, two target zones

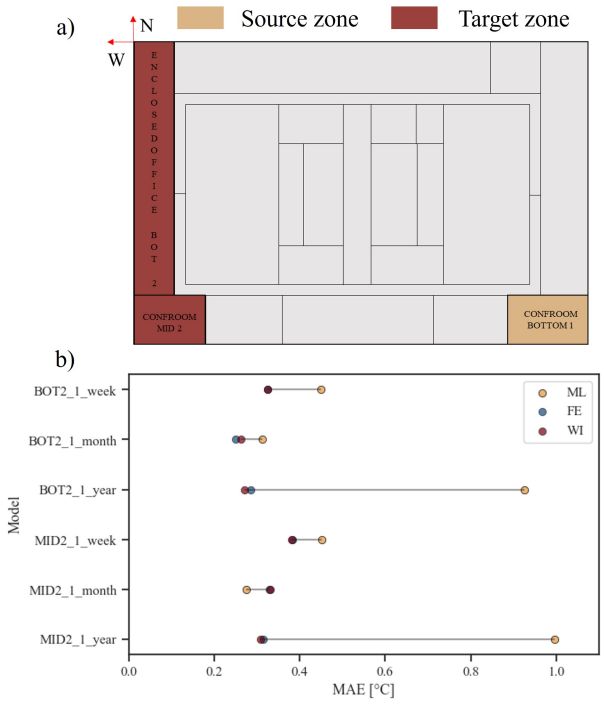

Figure 9: Performance of the different techniques over different zones

(highlighted in red in Figure 9a) were selected: a conference room on the second floor and an enclosed office on the second floor. The rationale behind zone selection was to test the neural networks with different orientation, area and floor. The conference room on the second floor (MID_2) was selected to test the influence of a different exposure on the model (changing it from east to west), while the enclosed office on the second floor (BOT_2) was selected to test both a different area and a different exposure. Once the zones were identified, the analysis was performed, considering different data availability (from one week to one year). Then several tests were performed that considered different data availability and compared the results of ML and TL models analysing the mean absolute error. Figure 9b shows that despite the different characteristics, TL was able to obtain better performance than standard ML independently from the data availability. Indeed, the ML model performance was heavily influenced by the amount of training data for the target building, while the TL model presented robust results over different data availability. After having assessed the ability of TL in different thermal zones, to isolate the effect of other variables, the following analysis was performed using the same thermal zone as a source.

Then, to analyse the average performance of the three techniques on the whole design of the experiments, mean absolute error was used to aggregate results over different climate, data availability, efficiency and occupancy profiles. As a result, Figure 10 shows the average MAE distribution for the three proposed approaches over all the simulations performed. The analysis of the distributions showed that ML trained over a period of one year in Climate 5A had in many cases unacceptable errors. A specific analysis will be conducted later to

understand the main factors related to the lower ML model performance. Furthermore, Figure 10a highlights the larger error distribution of the ML technique, which reaches values of more than 1 °C, while the TL maximum errors are below 0.7 °C. To better understand how the ML error is distributed, details for different data availability are shown in Figure 10b. The figure displays how one year of data led the ML model to a large error distribution, while one month of data showed the best performance, with an average error of 0.35 °C. As a result, the focus was shifted toward a one month training period. Figure 10c compares the error for each technique, assessing a slight performance improvement for both TL techniques over ML, with no particular differences between feature extraction and weight initialization.

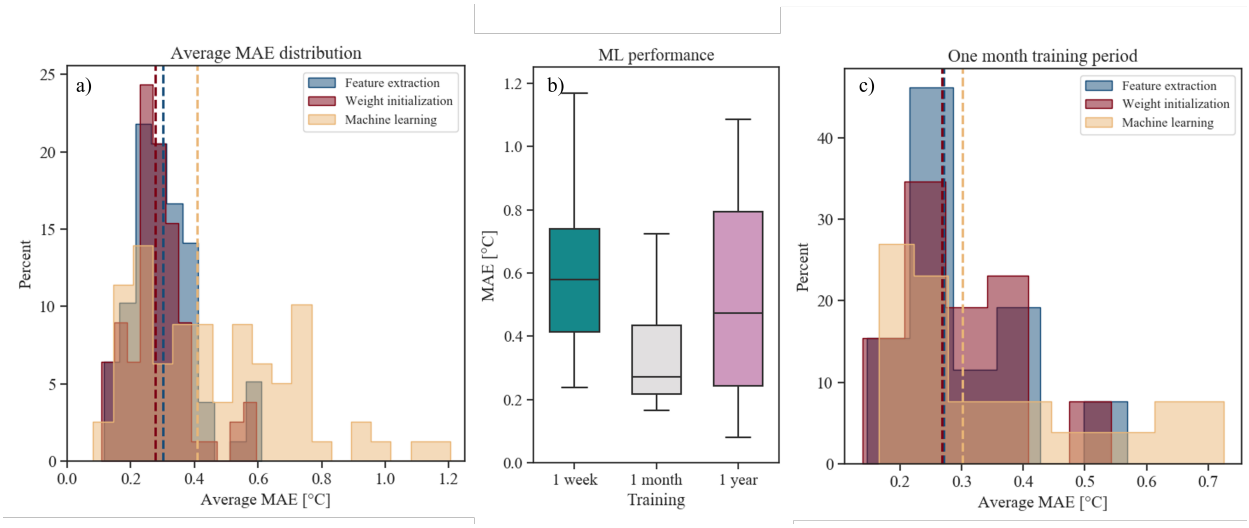

Figure 10: MAE distribution over different periods and techniques

To further study the effectiveness of transfer learning, average MAE distributions were divided in three ranges: low error (MAE < 0.4 °C), medium error (0.4 °C$<$ MAE < 0.7 °C), and high error (MAE > 0.7°C).

Figure 11 shows the error distribution by technique over all the influencing factors using a categorical plot. The ML technique is the only one with a high error, which mainly occurred with one week and one year of data. Furthermore, it shows how high errors are predominant in Climate 5A but are evenly distributed over the efficiency levels and occupancy runs. On the other hand, both feature extraction and weight initialization showed better performance; almost evenly distributed over different data availability, with lower error for Climate 3C, the same climate as that of the source building.

Due to the co-occurrence of different features on the model (e.g, different climate, occupancy and efficiency levels), a specific analysis was performed by changing only one feature at a time, with the goal of isolating their effect on model performance. Figure 12 shows the MAE for different techniques for several cases. Furthermore, it shows how by changing only the efficiency level (same climate and same occupancy profile), transfer learning outperforms machine learning for every data availability, while negative transfer can occur when buildings across different climates are analysed, with very different results according to

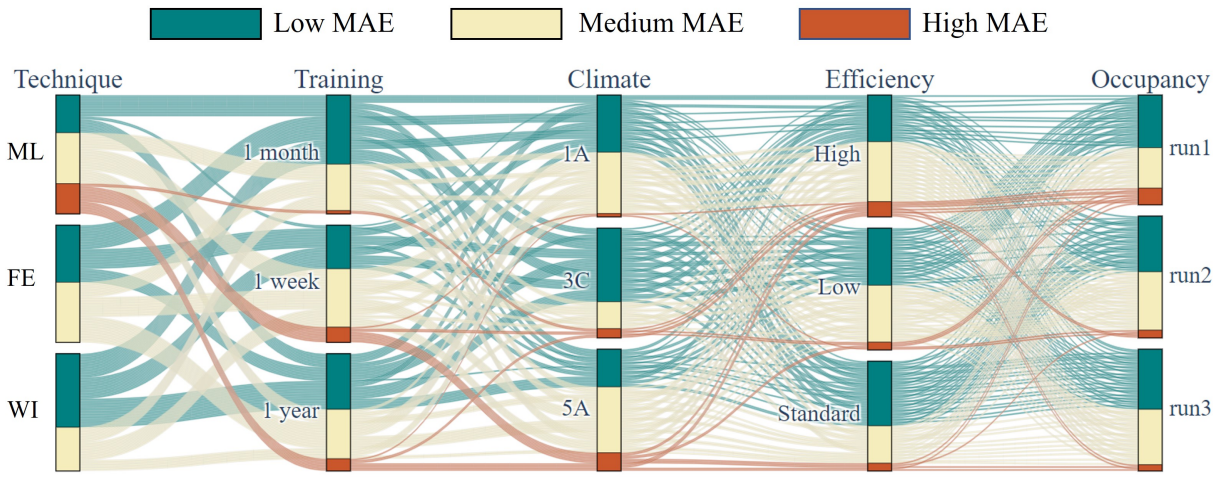

Figure 11: Methodological framework

data availability. Looking at results with various occupancy profiles, a narrow performance improvement can be seen, with a negligible case of negative transfer learning, since both ML and TL techniques have an average error below 0.2 °C and very similar performance.

Therefore, to assess the influence of climate and data availability on model performance, a specific analysis was conducted, as shown in Figure 13. In particular, Figure 13a shows the distribution of the mean absolute error for one week, one month and one year of data availability over the three different climates. For the sake of clarity, the error bar related to one year and Climate 5A, which exceeded 1.5 °C, has not been shown, while its lower outliers have been included in the figure. Note that often an MAE of 0.5 °C is seen as threshold for the deployment of a model that predicts the internal air temperature. As a result, the figure highlights the inadequacy of ML models to be deployed for the specific combination of climate and time horizon. With increasing data availability, the median value of the ML models decreases. In general, TL approaches are more robust compared to ML approaches. Furthermore, the analysis showed how almost every TL model had an error below 0.5 °C, while ML often exceeded this threshold. Figure 13b uses the asymptotic performance improvement to compare the simulation point by point. It can be seen how, on average, the best performance improvements are achieved in Climate 3C (i.e, the climate selected for the source building). Note that performance improvements for climate 5A are highly influenced by the poor performance of ML models, increasing the advantages of using TL. The main reason may be related to the high temperature variation of Climate 5A, which makes it hard for the model to generalize over the entire year. However, Figure 13b also highlights the presence of negative transfer, especially with one month of data, a period in which ML already has good performance. As a result, a further analysis was conducted to identify the main driver of negative TL.

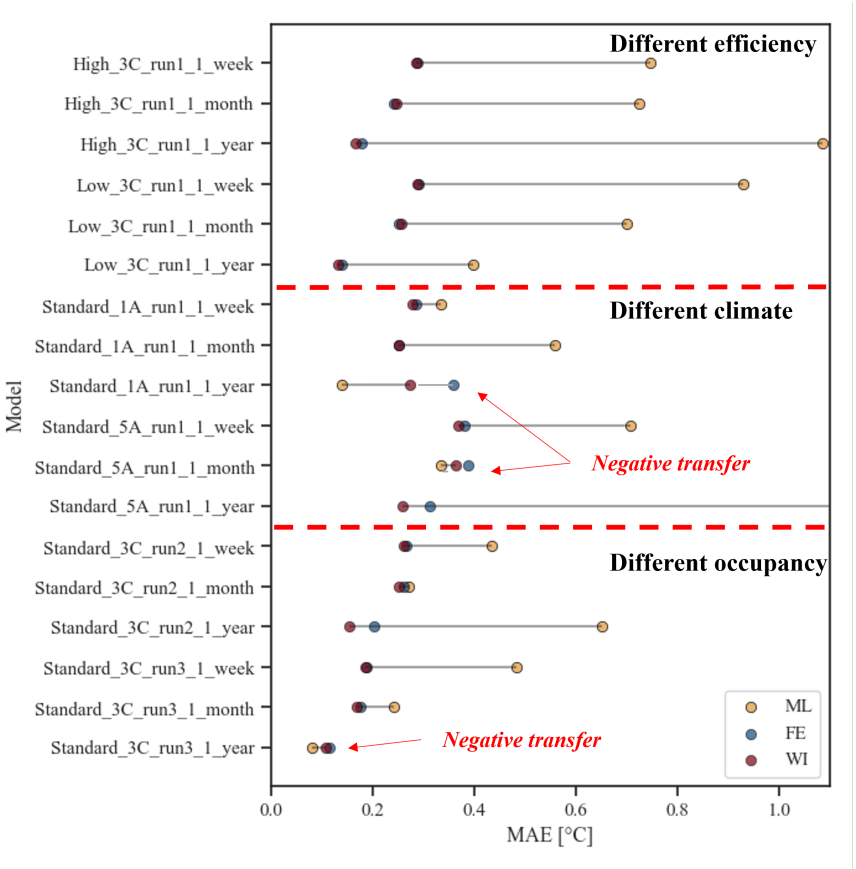

Figure 12: Performance comparison with isolated effects of features

5.2. Negative transfer learning

Figure 14 shows the asymptotic performance of all the simulations, highlighting three particular areas: negative transfer learning, neutral transfer and effective transfer. Negative transfer occurs when the MAE is greater than 0.05 °C compared to ML models, neutral transfer is when it is smaller than 0.1 °C, and effective transfer reduces the MAE at least 0.1°C. As can be seen, about 20% of cases have negative transfer, 20% have neutral transfer and 60% of cases show effective transfer. Figure 14b displays a detail of negative transfer, using different shapes and colors to highlight data availability and climate. The figure highlights how negative transfer occurred only 4 times out of 52 simulations, when one week of data was used (turquoise color), suggesting an effectiveness of TL in over 90% of the cases when one week of data is considered. It also can be noticed how negative transfer occurred only four times out of 52 simulations (diamond shape), with a performance increase in about 90% of the cases when the target building had the same climate as the source building. The figure shows that the highest amount of negative TL happened with one month of data availability, identifying this amount of data as enough to obtain a good ML model performance.

Lastly, to provide a comparison of the model performance with effective and negative

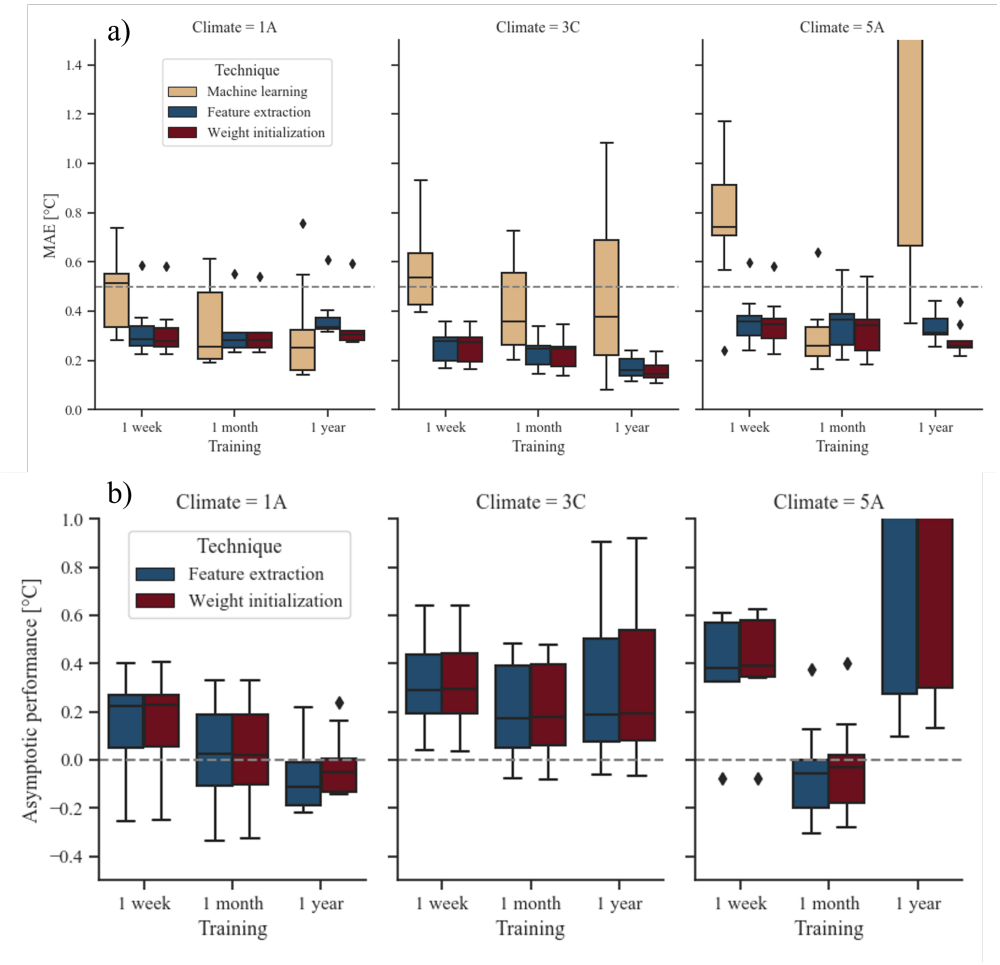

Figure 13: Methodological framework

transfer, Figure 15 displays temperature evolution for the first predicted time-step over a random day for real values using ML and TL models. The figure on the left highlights how in this case ML was not able to properly describe the building dynamics, while both TL techniques follow the trend of the real value (green). On the other hand, the right figure shows a case of a negative TL, in which the performance of TL was still able to capture the building dynamic but perform worse than the classical ML approach.

5.3. Jumpstart performance

Figure 16 shows jumpstart performance with different data availability. As shown, the highest jumpstart occurred for one month and one week, reducing the MAE of the first epoch about 8°C. Despite this reduction, the final performance during training was comparable to an ML model. Moreover, the figure shows how the performance of transfer learning is almost constant, thus highlighting the possibility of great computational cost reduction when using transfer learning. Transfer learning also provided a computational advantage; however, the model complexity and the time required to train such models in this kind of problem is little.

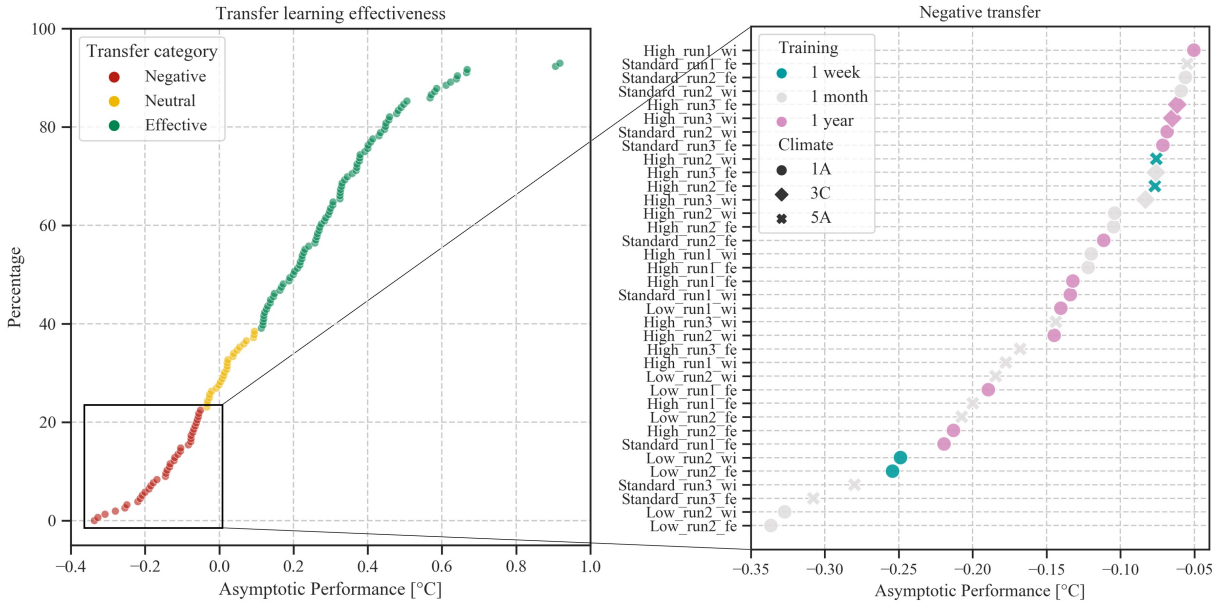

Figure 14: Categorization of transfer learning effectiveness and negative transfer analysis

These advantages are usually more important when dealing with different applications, such as in computer vision. As a result, the jumpstart performance is a less reliable metric compared to the asymptotic performance, which is better suited to quantify the goodness of a model.

5.4. Online deployment

Figure 17 shows the MAE error distribution over each week of the year for the two techniques (ML and TL) deployed online. The configuration selected for the target building was characterized by Climate 3C (the same of the source building), occupancy pattern 2 and a high efficiency level. This configuration was selected on the basis of the outcome of the previous analysis. The transfer learning model (already trained on one year of source building data) was updated for the target building each week as new data became available following an anchored deployment configuration that employed existing data and a new week's data using the same learning rate of transfer learning configuration ($2*10^{-3}$). The machine learning model used the same deployment strategy without leveraging pre-training data from the source building. To highlight both relative and absolute performance, the figure reports a candlestick visualization. The green color of the box highlights the cases when the TL showed higher performance against ML, while the red box represents the opposite occurrence. The height of the box measures the difference in terms of performance between the two models, while the two extremes indicate the absolute value of MAE. The figure shows that especially during the first weeks of deployment, the ML had very poor performance when compared with TL. However, as training data became available for the

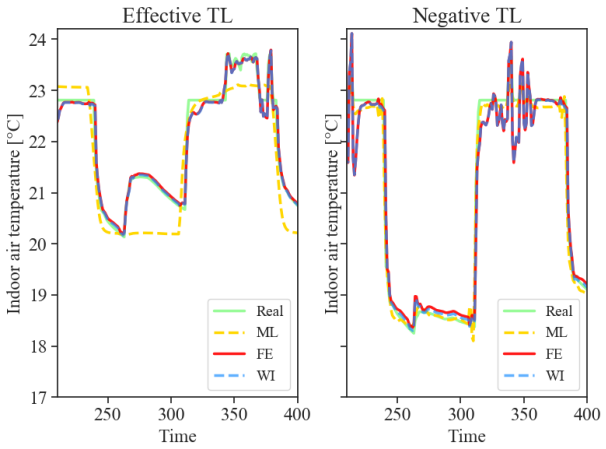

Figure 15: Prediction evolution for the first time-step with different techniques for effective and negative TL

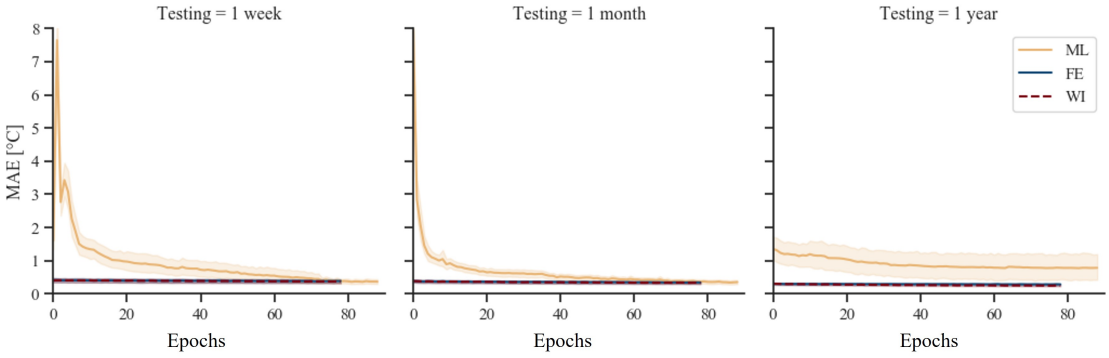

Figure 16: Jumpstart comparison over different training time

ML, the performance difference between the two models tended to decrease, and after week 40, the performance of the two models were comparable.

6. Discussion

Building dynamics prediction proved to be effective to unlock the potential of advanced control strategies. However, the main bottleneck is represented by the data availability in most of the buildings, making the exploitation of data driven models a niche. TL promises to overcome this problem, but still requires further studies to quantify building similarity. This research aimed to quantify the feature importance of several variables in a TL setting. In particular, this study compared two TL techniques and assessed the effect of data availability and case specific features (e.g., climate, efficiency level, occupancy). To capture the effect of the different variables, an experiment design was conceived. Analysis of the results revealed several insights: first, unlike the ML models, the error performance over multiple time-steps is very small for the first time-step, increasing more steeply with the following time-

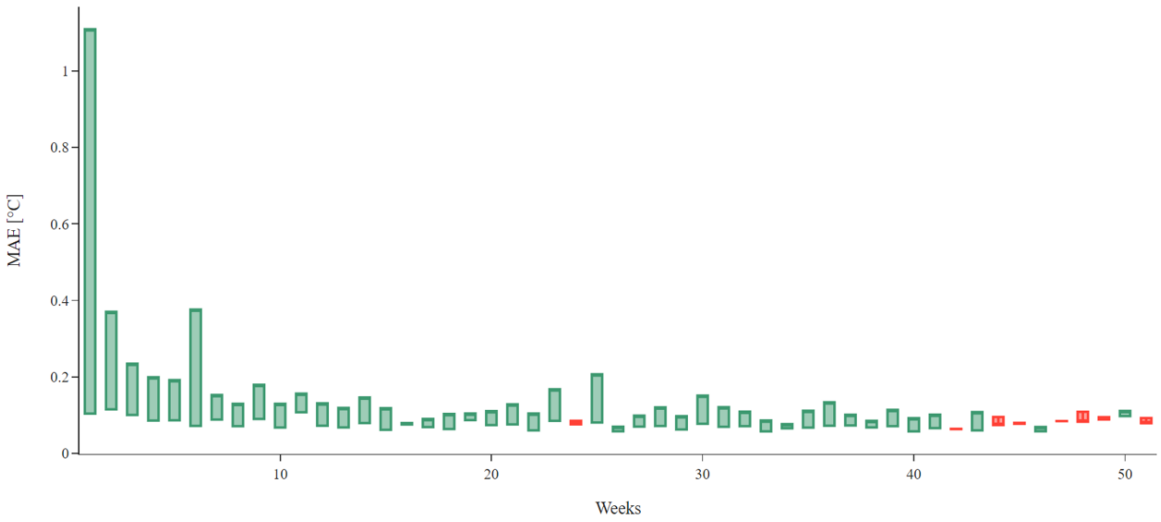

Figure 17: Performance comparison between online ML and online TL

steps, while on average TL showed better performance. This information is particularly helpful for advanced controllers that require precise estimation of the upcoming time-steps. Additionally, analysis of the data availability aimed to assess how much data are necessary in TL and ML settings. The analysis showed that for ML a higher amount of data in the target building may be counterproductive, especially when the target building is located in climates with a great variation between the different seasons (Climate 5A). On the other hand, using a large amount of data helps to reduce the variance of TL models, obtaining more robust results.

Furthermore, the analysis confirmed the ability of TL to deal with different efficiency levels and occupancy, while limitations were observed for its effective applications across different climates, highlighting the role of external (outdoor air) temperature as the most important feature. Moreover, the focus on the asymptotic performance and the negative transfer allowed researchers to identify guidelines and constraints on the application of transfer learning for building dynamics prediction. The analysis of the results identified data-scarcity (one week) and the application on the same climate of the source building as the best case study to deploy TL. Furthermore, performance analysis suggest that for different features (e.g., climate), when new data are available the optimal solution consists in using online transfer learning and shifting to online machine learning when a robust dataset in the target building is available (e.g., one year). On the other hand, when the most important features are the same for the source and target building, transfer learning may achieve performance improvement independently from the amount of data used, highlighting its effectiveness in generalizing machine learning models. Lastly, a specific analysis was carried out on jumpstart performance; however, despite the computational advantages introduced by TL, the time needed to train such models is relatively low, due to the small dimension with respect to computer vision domains. This suggests the use of asymptotic performance as a key performance indicator to evaluate the effectiveness of TL.

6.1. Limitations

A key concern about the comparison of the different models is the influence of the testing period on performance. Indeed, a model trained on one year of data and tested on one month may experience different performance based on the month selected for testing (e.g., winter, summer). This can be caused by considering the climate used for training, which can be cooling or heating dominant. However, the purpose of that type of analysis was outside the scope of this study and aimed at evaluating the robustness of ML and TL models. Furthermore, the influence of the thermal zone (e.g., a perimeter zone with more exposure to weather versus an interior zone) has not been fully characterized, since its introduction in the design of experiments would have reduced the output interpretability. Lastly, the limitation of the optimization process lies in the neural network structure, optimized only for the source case with a specific set of features (e.g., climate, efficiency level, occupancy). As a result, other case studies may have sub-optimal DNN architectures, especially for one year of data, that may require a more complex model. However, the methodology avoided the optimization of the DNN for every condition to allow a peer-to-peer comparison between ML and TL.

7. Conclusions

The present work introduced a methodology to assess feature importance in TL settings for building dynamics prediction. The problem involved the development of approximately 250 models that were deeply analysed in order to assess their most important features. The error distribution comparison showed that transfer learning reduced error in the first time-steps, suggesting its coupling with advanced control strategies, which rely on short-term predictions. The analysis also demonstrated that climate is the most influential feature, determining the success of the transfer. Another analysis showed that data availability also affects performance, suggesting the coupling of transfer learning with online learning if the features are extremely different, while showing that TL can adapt easily to different building energy efficiency levels and occupancy patterns. The study also compared two TL techniques (weight-initialization and feature extraction), showing a slightly better performance of weight-initialization. A specific analysis of negative transfer learning identified the strengths and weaknesses of the method. Lastly, this study compared the performances of the algorithm in an online fashion, to support its coupling with an advanced controller in real-world implementations. In conclusion, TL was found to be effective in most of the cases, especially when dealing with changes in building related variables, providing more robust predictions. Future studies will focus on:

- The application of transfer learning with different input data dimensions. Buildings have different levels of sensing and metering, providing diverse sources of measured data, which may represent a limit for the deployment of transfer learning models. Studies can aim to assess transfer learning's ability to adapt to different input dimensions without losing the ability to capture building thermal dynamics.

- The feasibility of developing a database of model archetypes for the most important features. This work aims to create different data-driven archetype models based on climate and building archetypes, which could be used by researchers to test and speed transfer learning applications in both commercial and residential buildings.
- The deployment of transfer learning techniques in residential buildings based on real data (e.g., the large scale smart thermostat data from ecobee's Donate Your Data program). Residential buildings have usually fewer sensors installed and are characterised by high stochasticity (due to occupant use behavior). The analysis can aim to assess the effectiveness of transfer learning on real data at a large scale.
- The implementation of transfer learning to simulate building thermal dynamics in a real building, to support control applications and evaluate its performance when integrated with advanced control strategies like model predictive control, comparing it against a traditional rule-based controller. The main goal of the application would be to study the interaction between neural network prediction and control strategies in a real building.
- The formulation of a transfer learning framework for control problem. The application would test the potentialities of transfer learning for advanced control (e.g., model predictive control, reinforcement learning) for performance benchmarking in a standardized test framework such as BOPTEST [42].

8. Acknowledgments

Li and Hong's work was supported by the Laboratory Directed Research and Development (LDRD) Program of Lawrence Berkeley National Labora- tory, provided by the Director, Office of Science, of the U.S. Department of Energy under Contract No. DE-AC02-05CH11231.

Nomenclature

ARMAX Autoregressive-Moving-Average with Exogenous Inputs

BNN Bayesian Neural Network

DOE Department of Energy

DRL Deep Reinforcement Learning

GEB Grid-interactive Efficient Building

HVAC Heating, Ventilation and Air Conditioning

LSTM Long Short Term Memory

MAE Mean Absolute Error

MAPE Mean Absolute Percentage Error

MEL Miscellaneous Electric Load

ML Machine Learning

MLP Multy-Layer Perceptron

MPC Model Predictive Control

MSE Mean Squared Error

PINN Physics-Informed Neural Network

RNN Recurrent Neural Network

TL Transfer Learning

VAV Variable Air Volume

XGBoost EXtreme Gradient Boosting

References

[1] M. Victoria, K. Zhu, T. Brown, G. B. Andresen, M. Greiner, Early decarbonisation of the European energy system pays off, Nature Communications 11 (2020) 1–9.

[2] A. Satchwell, M. A. Piette, A. Khandekar, J. Granderson, N. M. Frick, R. Hledik, A. Faruqui, L. Lam, S. Ross, J. Cohen, et al., A national roadmap for grid-interactive efficient buildings (2021).

[3] J. Drgoňa, J. Arroyo, I. Cupeiro Figueroa, D. Blum, K. Arendt, D. Kim, E. P. Ollé, J. Oravec, M. Wetter, D. L. Vrabie, L. Helsen, All you need to know about model predictive control for buildings, Annual Reviews in Control 50 (2020) 190–232.

[4] S. Aggarwal, R. Orvis, Grid flexibility: Methods for modernizing the power grid, Energy Innovation San Francisco, California March (2016).

[5] í Ciglera, D. Gyalistrasb, V.-N. Tietd, Luká, Ferkla, Beyond theory : the challenge of implementing model predictive control in buildings ji ř, 2013.

[6] H. Gao, C. Koch, Y. Wu, Building information modelling based building energy modelling: A review, Applied Energy 238 (2019) 320–343.

[7] D. Blum, K. Arendt, L. Rivalin, M. Piette, M. Wetter, C. Veje, Practical factors of envelope model setup and their effects on the performance of model predictive control for building heating, ventilating, and air conditioning systems, Applied Energy 236 (2019) 410–425.

[8] Z. Afroz, G. Shafiullah, T. Urmee, G. Higgins, Modeling techniques used in building hvac control systems: A review, Renewable and Sustainable Energy Reviews 83 (2018) 64–84.

[9] G. Pinto, M. S. Piscitelli, J. R. Vázquez-Canteli, Z. Nagy, A. Capozzoli, Coordinated energy management for a cluster of buildings through deep reinforcement learning, Energy 229 (2021) 120725.

[10] J. Drgoňa, D. Picard, M. Kvasnica, L. Helsen, Approximate model predictive building control via machine learning, Applied Energy 218 (2018) 199–216.

[11] A. Ruano, E. Crispim, E. Conceição, M. Lúcio, Prediction of building's temperature using neural networks models, Energy and Buildings 38 (2006) 682–694.

[12] C. Sun, J. Chen, S. Cao, X. Gao, G. Xia, C. Qi, X. Wu, A dynamic control strategy of district heating substations based on online prediction and indoor temperature feedback, Energy 235 (2021) 121228.
[13] X. Shi, W. Lu, Y. Zhao, P. Qin, Prediction of indoor temperature and relative humidity based on cloud database by using an improved bp neural network in chongqing, IEEE Access 6 (2018) 30559–30566.
[14] A. Kusiak, G. Xu, Modeling and optimization of hvac systems using a dynamic neural network, Energy 42 (2012) 241–250. 8th World Energy System Conference, WESC 2010.
[15] G. Mustafaraj, G. Lowry, J. Chen, Prediction of room temperature and relative humidity by autoregressive linear and nonlinear neural network models for an open office, Energy and Buildings 43 (2011) 1452–1460.
[16] Z. Afroz, T. Urmee, G. Shafiullah, G. Higgins, Real-time prediction model for indoor temperature in a commercial building, Applied Energy 231 (2018) 29–53.
[17] H. Huang, L. Chen, E. Hu, A neural network-based multi-zone modelling approach for predictive control system design in commercial buildings, Energy and Buildings 97 (2015) 86–97.
[18] Z. Wang, T. Hong, M. A. Piette, Building thermal load prediction through shallow machine learning and deep learning, Applied Energy 263 (2020) 114683.
[19] F. Mtibaa, K.-K. Nguyen, M. Azam, A. Papachristou, J.-S. Venne, M. Cheriet, Lstm-based indoor air temperature prediction framework for hvac systems in smart buildings, Neural Comput. Appl. 32 (2020) 17569–17585.
[20] C. Xu, H. Chen, J. Wang, Y. Guo, Y. Yuan, Improving prediction performance for indoor temperature in public buildings based on a novel deep learning method, Building and Environment 148 (2019) 128–135.
[21] M. J. Ellis, V. Chinde, An encoder–decoder lstm-based empc framework applied to a building hvac system, Chemical Engineering Research and Design 160 (2020) 508–520.
[22] Z. Fang, N. Crimier, L. Scanu, A. Midelet, A. Alyafi, B. Delinchant, Multi-zone indoor temperature prediction with lstm-based sequence to sequence model, Energy and Buildings 245 (2021) 111053.
[23] G. Pinto, D. Deltetto, A. Capozzoli, Data-driven district energy management with surrogate models and deep reinforcement learning, Applied Energy 304 (2021) 117642.
[24] M. Raissi, P. Perdikaris, G. Karniadakis, Physics-informed neural networks: A deep learning framework for solving forward and inverse problems involving nonlinear partial differential equations, Journal of Computational Physics 378 (2019) 686–707.
[25] F. Bünning, B. Huber, A. Schalbetter, A. Aboudonia, M. Hudoba de Badyn, P. Heer, R. S. Smith, J. Lygeros, Physics-informed linear regression is competitive with two machine learning methods in residential building mpc, Applied Energy 310 (2022) 118491.
[26] G. Gokhale, B. Claessens, C. Develder, Physics informed neural networks for control oriented thermal modeling of buildings, ArXiv abs/2111.12066 (2021).
[27] J. Drgoňa, A. R. Tuor, V. Chandan, D. L. Vrabie, Physics-constrained deep learning of multi-zone building thermal dynamics, Energy and Buildings 243 (2021) 110992.
[28] L. Di Natale, B. Svetozarevic, P. Heer, C. N. Jones, Physically consistent neural networks for building thermal modeling: theory and analysis, arXiv preprint arXiv:2112.03212 (2021).
[29] S. J. Pan, Q. Yang, A survey on transfer learning, IEEE Transactions on Knowledge and Data Engineering 22 (2010) 1345–1359.
[30] G. Pinto, Z. Wang, A. Roy, T. Hong, A. Capozzoli, Transfer learning for smart buildings: A critical review of algorithms, applications, and future perspectives, Advances in Applied Energy 5 (2022) 100084.
[31] M. M. Hossain, T. Zhang, O. Ardakanian, Evaluating the feasibility of reusing pre-trained thermal models in the residential sector, UrbSys 2019 - Proceedings of the 1st ACM International Workshop on Urban Building Energy Sensing, Controls, Big Data Analysis, and Visualization, Part of BuildSys 2019 (2019) 23–32.
[32] Z. Jiang, Y. M. Lee, Deep transfer learning for thermal dynamics modeling in smart buildings, in: 2019 IEEE International Conference on Big Data (Big Data), IEEE, 2019, pp. 2033–2037.
[33] Y. Chen, Y. Zheng, H. Samuelson, Fast adaptation of thermal dynamics model for predictive control

of hvac and natural ventilation using transfer learning with deep neural networks, in: 2020 American Control Conference (ACC), IEEE, 2020, pp. 2345–2350.
[34] Y. Chen, Z. Tong, Y. Zheng, H. Samuelson, L. Norford, Transfer learning with deep neural networks for model predictive control of hvac and natural ventilation in smart buildings, Journal of Cleaner Production 254 (2020) 119866.
[35] T. Grubinger, G. C. Chasparis, T. Natschläger, Generalized online transfer learning for climate control in residential buildings, Energy and Buildings 139 (2017) 63–71.
[36] C. Tan, F. Sun, T. Kong, W. Zhang, C. Yang, C. Liu, A survey on deep transfer learning, Lecture Notes in Computer Science (including subseries Lecture Notes in Artificial Intelligence and Lecture Notes in Bioinformatics) 11141 LNCS (2018) 270–279.
[37] S. Hochreiter, J. Schmidhuber, Long short-term memory, Neural Computation 9 (1997) 1735–1780.
[38] Department of Energy, Commercial reference buildings, 2022. URL: https://www.energy.gov/eere/buildings/commercial-reference-buildings.
[39] Y. Chen, T. Hong, X. Luo, An agent-based stochastic occupancy simulator, Building Simulation 11 (2017) 37–49.
[40] H. Li, Z. Wang, T. Hong, A synthetic building operation dataset, Scientific Data 8 (2021) 213.
[41] T. Akiba, S. Sano, T. Yanase, T. Ohta, M. Koyama, Optuna: A next-generation hyperparameter optimization framework, 2019. URL: https://arxiv.org/abs/1907.10902. doi:10.48550/ARXIV.1907.10902.
[42] D. Blum, J. Arroyo, S. Huang, J. Drgoňa, et al., Building optimization testing framework (boptest) for simulation-based benchmarking of control strategies in buildings, Building Performance Simulation 14 (2021) 586–610.