SimBuild 2024

Eleventh National Conference of IBPSA-USA
Denver, Colorado
May 21 – 23, 2024

A Parameter-based Transfer Learning Approach for Predicting Occupancy in Institutional Buildings

Aya Doma[1], Fatima Amara[2], and Mohamed Ouf[1]
[1]Department of Building, Civil, and Environmental Engineering, Concordia University, Canada
[2]Hydro Quebec, Quebec, Canada

Abstract

Accurately representing buildings' occupancy schedules has been crucial in understanding the energy flexibility of buildings as well as improving the allocation of building services and resources. However, the limited availability of data to model occupancy schedules for specific buildings has been challenging the development of an accurate prediction model. This study evaluates a parameter-based transfer learning scheme developed for time series prediction of occupancy schedules with limited data. The study focuses on an institutional building in Quebec, Canada as a case study. The results showed that transferring the knowledge of the pre-trained time series model has resulted not only in a reduction in the computational power allocated to the fitting of the model but also reduced the model's training requirements while maintaining a mean absolute error of 3 occupants.

Introduction

Building occupancy information and schedules have been acknowledged as a critical factor for realistically planning the ongoing transition to a zero-emission grid and the electrification of the different sectors. On a building scale, previous studies that focused on the automation of building systems control and operation emphasized the effect of considering occupancy schedules and levels of variation to achieve more efficient energy consumption (Drgoňa et al. 2020; Mirakhorli and Dong 2016; Yang et al. 2022). Sharma et al. (2022) showed that using occupancy information in model predictive control of an office building can achieve a significant reduction in building energy consumption without sacrificing occupant comfort. On an urban scale, research efforts focused on accommodating the urgent need for flexible load controls to ensure grid reliability through different demand-side management strategies. They found that occupancy-based buildings-to-grid integration strategies are more sustainable and can provide up to 60% of potential energy cost savings (Dong et al. 2018). The occupancy models in this case are mainly used to quantify the energy flexibility of each building with respect to the occupancy levels and schedules and accordingly, identify the periods in which the building can contribute and participate in the different demand and supply management programs. In addition to these roles, realistic and accurate occupancy schedules as well as occupant behavior patterns have been widely recognized as one of the main sources of the discrepancies between building energy simulation results and the actual energy use in buildings on both building and urban scales (Doma and Ouf 2023; Hong, Langevin, and Sun 2018; Dabirian, Panchabikesan, and Eicker 2022). Accordingly, analyzing occupancy patterns for the different building sectors and developing a reliable prediction model for future planning has been an interest of many research efforts. They mainly investigated the use of different technologies to collect representative occupancy information and the different modelling approaches for accurately representing the stochastic nature of occupancy.

In terms of data collection, many researchers relied on occupancy information sensing technologies that can explicitly report occupancy status or count. The most commonly used method is the passive infrared (PIR) sensors where occupants' movements are detected and the room occupancy status is identified accordingly (Shen, Newsham, and Gunay 2017). Estimating the number of occupants for the whole building using CO2 readings has also attracted the interest of researchers (Arora et al. 2015; Dong and Andrews 2009; Tekler and Chong 2022). However, the readings from these sensors are usually coupled with other measurements for the indoor condition for more accurate occupancy estimation. Recently a few studies have relied on camera-based occupancy counters to detect and estimate the whole building occupant count (Zhao, Tu, and Chang 2019; Guan and Huang 2015). Regardless of the data collection method, occupancy patterns are not only expected to vary from one building to the other but also the pattern of one building is subjected to change over time due to transformations in space utilization patterns, technological advances, and the expected increase in telecommuting (Alishahi 2022; Hobson et al. 2020; Ouf, O'Brien, and Gunay 2019). Accordingly, even considering these counters' reliability to provide accurate occupancy estimation, limited data will be available with them being newly introduced to existing buildings. Considering data from the available infrastructure in buildings such as the Wi-Fi connection count can also face the same challenge of data inadequacy if the data was not collected before the beginning of the analysis (Shen, Newsham, and Gunay 2017). Moreover, obtaining such data from these sources at scale can face multiple logistical issues as well as privacy concerns.

Two main approaches have been investigated to address these challenges. The first is related to the data sources themselves while the other is related to the modelling techniques. In terms of the data sources, the mobile positioning dataset showed promising results when used to generate building occupancy profiles (Wu et al. 2020; Happle, Fonseca, and Schlueter 2020; Doma et al. 2022). These datasets are GPS-based allowing for tracking the movement of devices (i.e., smartphones) and then converting them into building occupancy patterns. However, these patterns are limited to specific points of interest (POIs). On the other hand, the transfer learning (TL) technique has been used to address the limited data size by reusing the pre-trained occupancy models (Park et al. 2022). This technique is based on the idea that knowledge gained while solving one problem can be applied to a different but related problem (Weiss, Khoshgoftaar, and Wang 2016). Accordingly, a developed machine learning model for a specific task can be reused as a starting point to provide learned knowledge and experience to a model on a second task. The TL technique has proven its efficiency in training prediction models using small datasets but only for neural network (NN) models. These models are more complicated to deploy and integrate with existing control algorithms in buildings. Recently, however, a modified transfer learning scheme was introduced for transferring pre-trained model knowledge for regression time series models (Warushavithana et al. 2021). The scheme was validated using simulation data from different domains, more specifically to predict indoor air quality over time, COVID-19 spread, and outdoor weather. However, to the authors' knowledge, the scheme has not yet been validated for predicting human behaviour such as the occupancy pattern with ground truth data.

To this end, this paper investigates the applicability of combining the two mentioned approaches for addressing data insufficiency by using transfer learning approaches to transfer the structure and the tuned parameters of a pre-trained time-series occupancy estimation model. The study focuses on institutional buildings, specifically public libraries in Quebec, Canada as a case study. The objectives of the study are as follows: i) leverage rich mobile positioning data to develop a time series forecasting model to predict day-ahead occupancy schedule, ii) analyze the effect of using building characteristics such as building opening hours as well as the outdoor environmental conditions as variables on improving model performance, and iii) analyze the effect of reusing the pre-trained model structure and parameters to initialize the fitting of the model to a target building with limited training data.

The next sections will start with covering the literature related to this study and giving a brief background, then the methodology followed to apply the objectives is discussed in detail followed by the results, and finally the conclusion.

Background

Transfer learning (TL) can be categorized based on the strategy adopted to share the knowledge into four main categories (Pinto et al. 2022). First, the instance-based TL with the main goal of incorporating data points from a source model into a target model to improve the prediction. The second category is the feature representation-based TL which aims to map instances between the source and the target models. On the other hand, the third category of the TL techniques follows the assumption that the data has similar relations in both the source and the target models, and accordingly, the knowledge related to the relationships between the two datasets is transferred. Finally, the parameter-based TL shares the parameters or the distributions of the hyper-parameters of the source model to accelerate the learning process of the target model with limited data.

Applying transfer learning techniques to the occupancy prediction model can be considered one of the relatively new areas for research interests. Accordingly, a limited number of models and algorithms were investigated. Khalil et al. (2021) developed a stacked long-short-term memory (LSTM) model to predict the occupancy status of an office building then relied on transfer learning techniques for neural networks to transfer the model into another two offices at the same building. Stjelja et al. (2022) used a similar approach but was applied to meeting rooms in hospitals instead of offices. Dridi, Amayri, and Bouguila (2022) on the other hand, investigated feature representation-based transfer learning algorithms to map the features of two datasets with the goal of estimating the occupancy status (no occupants, one occupant, or two occupants) of two offices. While the developed models from the mentioned studies focused on predicting the occupancy status instead of accurate occupancy number estimation, they also focused on testing the applicability of TL in predicting occupancy on a room or a zoon scale instead of investigating the applicability of transferring pre-trained occupancy model from one building to another.

Methodology

Figure 1 shows the overall methodology followed in this study. In the next subsections, the pre-trained model is explained in detail considering the data used to train the model, the source building where the data is collected, and the model development and performance evaluation steps. These subsections also include an explanation of the available data for the target building and the steps followed to transfer the information of the pre-trained model from the source to the target.

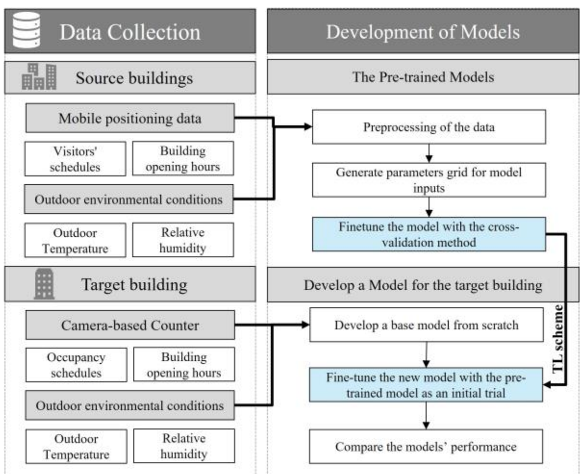

Figure 1: The overall research methodology.

The target and source buildings

The study focuses on developing an occupancy model for a public library located in Quebec, Canada (ASHRAE climate zone 6). This library will be referred to as the target building, or library 1 in this study. The target building consists of two main floors with a total net floor area of 2100 m^2. The two floors are mainly open-plan reading areas and bookshelves with a few offices on the side of the second floor. The building has one entrance/exit connected to the reception area as well as the stairs heading to the second floor.

The source building refers to the building with a large data size used to train and extract the occupancy model parameters. The chosen source building, which will be referred to as Library 2, is also located in the same Canadian province (i.e., Quebec) and the same climate zone. Library 2 is only one floor with an area of 1350 m^2. The library interior is an open plan of reading areas and bookshelves similar to the target building.

Data description

The target building has a camera-based occupancy counter installed at the building's only entrance. The sensor reports the number of people going in and out of the library each minute, then these values were used to calculate the occupancy inside the library. The data from the sensor was collected starting from April 15th, 2023, till October 24th,2023 (6 months of hourly data). Since the data size will not allow for examining the annual seasonality of the occupancy, data from a source library was chosen to train the pre-trained occupancy model.

The data from the source building (i.e., library 2) is the weekly patterns by SafeGraph.Inc (SG) dataset (SafeGraph Inc. 2023). The weekly pattern dataset is rich mobile positioning data collected from mobile devices opted into tracking their location. The main features and attributes in the data are presented in Table 1. As seen in the table, the data includes the foot traffic and visit schedules aggregated by point of interest (POI) to protect users' privacy. This data includes the hourly visit records for each POI from December 2018 to October 2022. However, data records from March 2020 to June 2021 were excluded from the analysis as they were during the COVID-19 pandemic which might cause the data to be biased.

The main steps for the preprocessing of the visitors' data to generate hourly occupancy schedules are explained in detail in (Doma et al. 2022). Figure 2 compares the average weekly patterns of both libraries normalized based on the peak occupancy for each building, while Table 2 compares the descriptive analysis of the occupancy of both buildings.

Table 1: SG features related to this study.

Attribute	Description
Placekey	Unique ID for each POI.
Coordinate	The latitude and longitude coordination of each POI.
wkt_area_sq_meters	The total area of the POI in square meters.
Visits by each hour	The number of visits to the POI for each of the 168 hours of the week.
POI opening hours	The opening duration for each day of the week for the POI.

Table 2: Descriptive analysis of the two libraries.

Descriptive Analysis	Target building: Library 1	Source building: Library 2
Floor Area (m2)	2100	1350
Data size	6 months	3 years
Data resolution	Hourly	Hourly
Average Occupancy	28	20

Maximum hourly occupancy	233	213

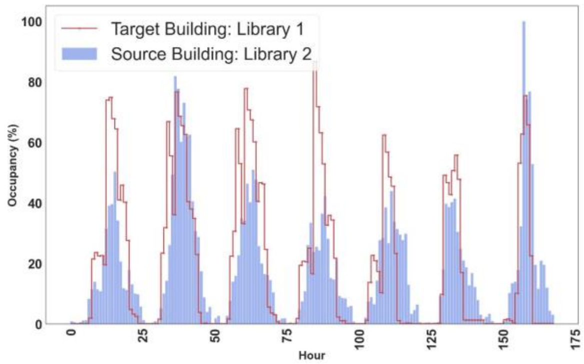

Figure 2: Comparing the normalized typical week of the target and the source buildings.

Time series forecasting models.

The occupancy prediction model should be able to capture the occupancy patterns of the building with respect to temporal, environmental, or spatial factors as well as model the seasonality of the patterns if needed. Moreover, as with any human activity, occupancy is stochastic, and accordingly, its randomness should be considered in the modelling process. Considering these requirements, time series forecasting techniques were chosen to model the occupancy in this study. In these forecasting techniques, the data is analyzed with respect to its chronological order, temporal factors, and the underlying correlations between the present and the historical data. Most of the widely used time series forecasting methods, such as the Autoregressive integrated moving average (ARIMA) family, assume the correlations between the target series and the regressors are linear, and require several preprocessing steps to configure the model components (Jin et al. 2021; Qiao, Yunusa-Kaltungo, and Edwards 2020). For the nonlinear correlations, models based on neural networks were commonly used, however, these models' interpretability of the correlation between the predictors and the targeted series is relatively low and is more complicated to deploy and integrate with the existing building operation and control systems.

Recently, Facebook introduced its time series forecasting model the Facebook Prophet (FBP) (Taylor and Letham 2018). The model was designed to capture linear and nonlinear patterns with relatively fewer preprocessing steps. Moreover, the model proved its ability to capture various patterns with minimal inputs of time series data which simplifies its integration into the control algorithm within the building automation systems (BAS). The model has been used for predicting hourly heating energy consumption (Jang, Han, and Leigh 2022), hourly indoor temperature (Fang et al. 2021), domestic hot water use schedule (Ivanko, Sørensen, and Nord 2020), as well as binary occupancy states based on readings from indoor environmental conditions sensors (Parise et al. 2021; Dorokhova, Ballif, and Wyrsch 2020). However, to the authors' knowledge, the prophet model has not been tested yet with predicting the occupancy count with different predictors and structures. Nevertheless, the ability of the FBP model to capture patterns and seasonality was proven as well as the model's simple integration with different control algorithms (Dorokhova, Ballif, and Wyrsch 2020).

The Facebook Prophet model.

The FBP model developed in this study is the generalized additive regression (AR) with two main layers: the components layer and the weight vector layer as seen in Figure 3. The former layer represents the growth (i.e., trend and pattern), the regular seasonality (yearly, weekly, and daily), and any conditional seasonality (e.g., holidays), while in the latter layer, the impact of these components is scaled with weight factors as well as an extra factor for the uncertainty in the pattern modelled. The growth of the data can be modelled as linear, logistic with specified capacity (i.e., maximum, and minimum growth), or flat based on observation from the data. Moreover, it is also represented in the model using changepoint detection analysis with the number of change points, their location, and their scale (i.e., the weight of their effect). While the first two factors should be extracted fully from the data, the scale of the change point is a parameter that should be tuned based on the characteristics of the modelled series (i.e., the occupancy). In the FBP algorithm, this scale is represented by a number to modulate the degree of the pattern flexibility to change according to the identified points. The large number represents more stochasticity in the modelled patterns where the effect of historical data is not always forced on the future prediction and different scenarios are explored.

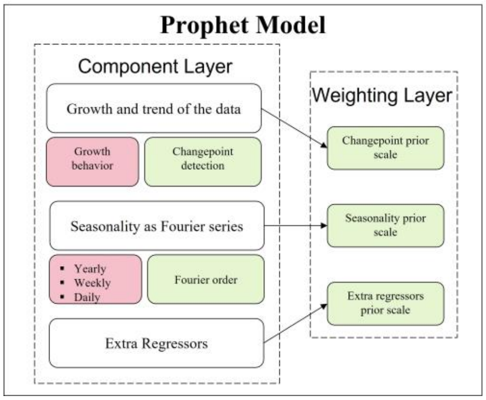

Figure 3: The FBP model structure and components.

The seasonality in the modelled pattern is examined with four different types: annual, weekly, daily, and temporary (conditional) seasonality. The first three types are modeled using the Fourier series and accordingly, they are represented with two main parameters: the scale and the Fourier order. The former parameter, similar to the changepoint scale, represents the degree of flexibility in applying the seasonality correlation extracted from the data, while the latter parameter describes how quickly the pattern changes due to seasonality. The model also allows examining the effect of temporary conditions (e.g., holidays) on the modelled pattern. To add any temporary condition, the starting and ending dates of the conditions are required. Considering these three components and their scaling factors, the FBP model can be presented as shown in equation (1).

$$Y_t = g_t + S_t + h_t + r_t + \varepsilon_t \qquad (1)$$

Where Y_t is the prediction of the model, in this study it is the occupancy number estimated at time t, g_t is the growth and trend component, S_t is the seasonality component, h_t is an optional component to represent the holidays, and r_t is another optional components in case of other time series variables or event should be considered (e.g., changing in the outdoor conditions). While the model learns the correlation between these components and the occupancy prediction, it also identifies the noise and the outliers (ε_t).

Development of the pre-trained model

In time series forecasting techniques and for the FBP model specifically, it is recommended to use one full year of data to better capture the growth and seasonality of the modelled pattern. However, since the camera-based counter was newly installed at Library 1, this requirement is not satisfied. Accordingly, a pre-trained FBP model was developed using the data from Library 2 to overcome this challenge. By fitting the FBP mode to the SG data from Library 2, it was found that the linear trend is the best representation of the occupancy growth at the library. Moreover, the configuration of the pre-trained model weight vectors and parameters is based on a grid search and the K-fold cross-validation (CV) method for time series data, as shown in Figure 4. At the end of the process, the hyperparameters of the models were selected. It should be mentioned that the choice of the ranges in the grid parameters used in this study was based on allowing the model to examine weak to strong changepoint and seasonality scales as well as different Fourier orders to represent fast and slow change in the patterns (Taylor and Letham 2018). Moreover, the scale and the Fourier order parameters were tested for each type of seasonality (yearly, weekly, and daily) separately.

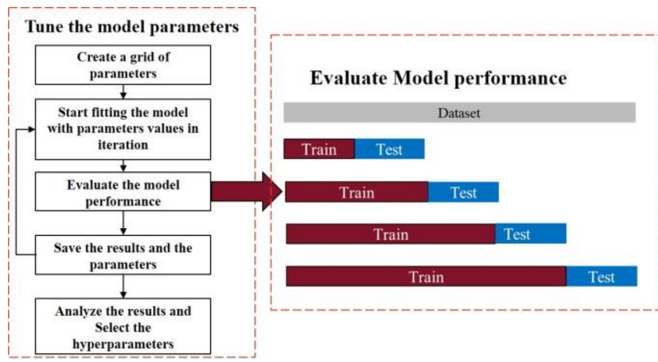

Figure 4: The cross-validation method to tune model parameters and evaluate its performance.

The performance of the developed model was evaluated using the mean absolute error (MAE) shown in Equation 2.

$$MAE = \frac{1}{n}\sum_{t=1}^{n}|Y_t - A_t| \qquad (2)$$

Where Y_t is the forecasted value as previously mentioned, A_t is the actual occupancy at time t, and n is the total number of points within the testing set.

Three different models are explored with the FBP algorithm. The three models are summarized and described in Table 3. The first model relies only on temporal variables extracted from the time series to predict the occupancy. In the second model, however, the impact of adding the opening hours of the libraries is examined by adding the opening hour regressor. The regressor is represented with a categorical variable to represent the status of the building at each timestep (closed, opened for employees only, open for visitors). Finally, the third model adds the outdoor environmental conditions as extra regressors to predict the occupancy count at each time step. The outdoor conditions were extracted from historical data based on the coordinates of the libraries and considered the outdoor temperature, wind speed, and relative humidity. These extra regressors (i.e., the opening hours and the outdoor conditions) are added using the component r_t.

Table 3: Description of the three pre-trained FBP models developed for library 2.

Model	Inputs
Model (1)	Time from the time series.
Model (2)	Time and the operational status are based on opening hours.
Model (3)	Time, operational status, and outdoor conditions.

Transferring the pre-trained model to the target building

In the parameters-based TL, the information collected from the source model is sent to another model that has the same task (predicting occupancy in this paper) in the form of shared model weights or for initialization of the fine-tuning process. A recent study developed a parameter-based transfer learning scheme to reuse the pre-trained Facebook prophet model with the goal of reducing the model requirements when trained on a new dataset (Warushavithana et al. 2021). This scheme is summarized and illustrated in Figure 5. The concept is based on training a FBP model exhaustively on a rich dataset, where the hyperparameters of the model are tuned using a grid search, and then the tuned parameters as well as using the weight vector layer can be used as initial starting points for optimizing the model on the new dataset. This process should help reduce the computational expenses of fitting and tuning the model on new data as well as reduce the data size requirements of the model to learn from the new data.

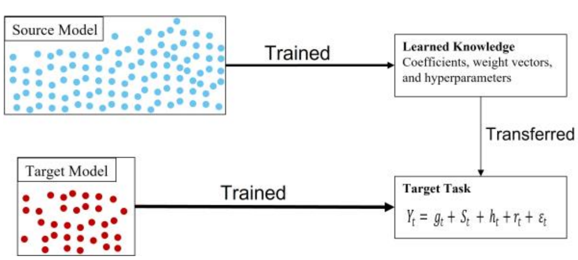

Figure 5: Transfer learning scheme for the FBP model.

Accordingly, the knowledge of the FBP model pre-trained with library 2 data, represented in the model components coefficients explained in the previous section, and the weight layer structure and values will be used as an initial value to fit the model to data from library 1. The fitting of the FBP model is based on the second optimization function of the Broyden, Fletcher, Goldfarb, and Shanno (BFGS) Algorithm (Taylor and Letham 2018). Using the knowledge of the pre-trained model as an initial value in the BFGS algorithm will determine both the direction and the step size to move in order to change the input parameters to minimize the objective function of the optimization, in this case, the error between the model prediction and actual values. To compare the impact of the process on the performance of the model, an independent model was developed to be used as a base model where the sensor data were used to train the model from scratch. Table 4 summarizes the models developed for library 1.

Table 4: Description of the FBP models developed for library 1.

Model	Parameters& structure	Inputs
Independent Model	Trained from scratch with library 1 data.	Time, operational status, and outdoor conditions
Model (1)	Transferred from Library 2	Time from the time series.
Model (2)	Transferred from Library 2	Time and the operational status are based on opening hours.
Model (3)	Transferred from Library 2	Time, operational status, and outdoor conditions.

Results and Discussion

As mentioned in the previous section, three models were developed using SG data from library 2. The first relied only on the temporal variables extracted from the timestamps of the time series, while the opening hours of the buildings as well as the outdoor conditions were added as extra regressors for the second and third models respectively. Figure 6 compares the performance of the three models when they were trained with different data sizes. As seen in the figure, the performance and accuracy of the models improved significantly when they were trained on two years of occupancy patterns in comparison to only one year. This highlights the importance and the impact of the data size on the accuracy of the model prediction.

In terms of comparing the performance of the three models, the figure shows that model 3, trained with outdoor conditions, performed slightly better than models 1 and 2 with a MAE of 7 occupants. The weight vector layer in the model ranked the outdoor conditions with the highest weight and effect on the occupancy pattern of the building followed by the changepoint detecting the daily peaks, while the weekly and annual seasonality were equally weighted and ranked as the least influential factors. While the impact of outdoor conditions is not fully shown with the model trained with 12 months of dataset, it can be observed when the training dataset expanded to include 24 months of data. As seen in the figure, with 12 months of training dataset, model 1 performed slightly better, on average, than model 3. On the contrary, this trend is reversed with an expanded training dataset of 24 months, where model 3 performed significantly better, on average, than model 1. According to these findings, it can be argued that the inclusion of outdoor conditions as predictors can boost the model performance, but this improvement becomes apparent only when the model is trained with an extensive, rich dataset collected over a prolonged period.

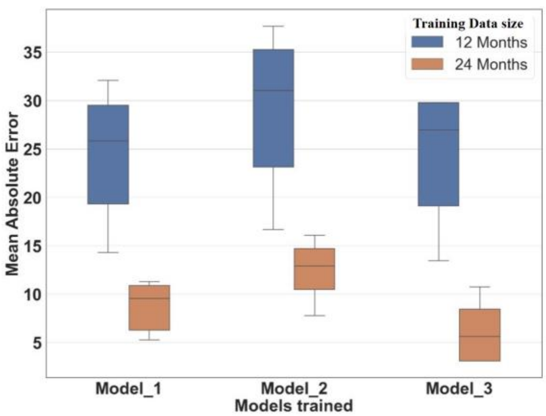

Figure 6: Comparing the performance of the three models developed for library 2.

Figure 7 shows the results of the models developed for Library 1 and compares the results of the base models with the three transferred models. As seen from the figure, the independent model trained from scratch reported a higher MAE of 5±2 occupants compared to the transferred models. Moreover, the performance of the model did not improve by much when more data records were added. This shows that the model was not able to capture the correlation between these regressors and the occupancy patterns using a small dataset. On the other hand, the transferred models showed faster learning ability than the one trained from scratch which can be seen by improving the models' performance with each month added to the training dataset. With four months of data fitted into the transferred models, significant improvements in the performance and the accuracy of the models' estimation were seen in comparison with the model trained from scratch. The MAE of the transferred models improved by almost 50% with the lowest error of 2 occupants using the opening hours of the library as a regressor. It is also important to mention that the outdoor conditions did not have much impact on the performance of the models developed for Library 1 compared to the impact they showed in the developed models for Library 2. This expected observation indicated that showing the variety and the changing of the outdoor conditions with more data records will help capture the correlation between them and occupancy patterns.

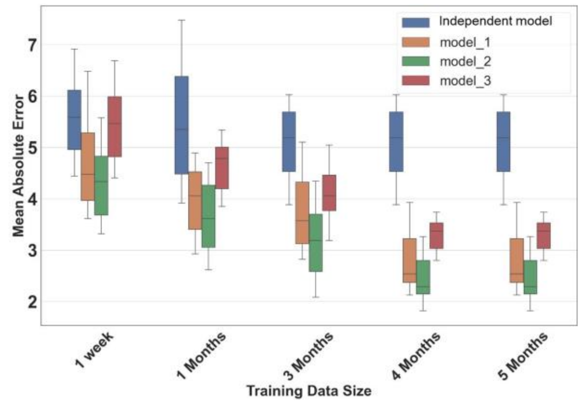

Figure 7: Comparing the performance of the models developed for library 1.

Conclusion

This paper evaluated the applicability of reusing and transferring the knowledge of a pre-trained time series model that predicts a day-ahead hourly occupancy schedule for an institutional building to overcome the challenges of data availability when training occupancy models from scratch. In addition, this paper also investigated the impact of using building operation hours as well as outdoor conditions on improving the accuracy of occupancy prediction models. The Facebook prophet model was used to predict the hourly occupancy of a library building in Quebec using three years of data for training. The results showed that training the model using the outdoor conditions as extra predictors next to the temporal variables improved the accuracy of the model and reduced the mean absolute error of the model down to 7 occupants (3%) compared to 12 occupants (5%). The knowledge of the model, which is represented in the model coefficient, weight vectors, and hyperparameters was then transferred and used to initiate the fitting of the model to a new dataset from a different building. The data from the target building was collected from a camera-based occupancy counter for six months.

The results showed that by transferring the knowledge of a pre-trained model the model requirements related to the training data size can be reduced while improving the accuracy of the model. Future works will focus on expanding and improving the transfer learning for the time series scheme to develop a generalized occupancy prediction model for different building sectors with the goal of providing the necessary occupancy inputs for urban building energy simulation. Furthermore, investigating the role and applications of the developed occupancy models within building-to-grid integration programs will be another interest of future research. As buildings' role in the grid is now transferring to prosumers who are affecting the supply and demand side, developing occupancy prediction models for the different buildings can play a significant role in transactive energy markets and managing the demand-side programs. Therefore, future work will investigate the deployment of the developed occupancy models into the current and future electric power systems and grid infrastructures.

Acknowledgment

References

Alishahi, Nastaran. 2022. "Using WiFi Connection Counts and Camera-Based Occupancy Counts to Estimate and Predict Building Occupancy."

Arora, Abhay, Manar Amayri, Venkataramana Badarla, Stéphane Ploix, and Sanghamitra Bandyopadhyay. 2015. "OCCUPANCY ESTIMATION USING NON INTRUSIVE SENSORS IN ENERGY EFFICIENT BUILDINGS." In 14th Conference of International Building Performance Simulation Association, Hyderabad, India, Dec. 7-9, 2015. Proceedings of BS2015. Hyderabad, India. https://hal.science/hal-01864941.

Dabirian, Sanam, Karthik Panchabikesan, and Ursula Eicker. 2022. "Occupant-Centric Urban Building Energy Modeling: Approaches, Inputs, and Data Sources - A Review." Energy and Buildings 257 (February): 111809. https://doi.org/10.1016/j.enbuild.2021.111809.

Doma, Aya, Mohamed Osman, Ahmed Zaalouk, Hamidreza Karimian, Mazdak Nik-Bakht, and Mohamed Ouf. 2022. "Towards Post-Pandemic Occupancy Patterns: Investigating Building Occupancy in New York City." ASHRAE Transactions 128 (1): 323–30.

Doma, Aya, and Mohamed Ouf. 2023. "Modelling Occupant Behaviour for Urban Scale Simulation: Review of Available Approaches and Tools." Building Simulation 16 (2): 169–84. https://doi.org/10.1007/s12273-022-0939-3.

Dong, Bing, and Burton Andrews. 2009. "Sensor-Based Occupancy Behavioral Pattern Recognition for Energy and Comfort Management in Intelligent Buildings," June. https://www.aivc.org/resource/sensor-based-occupancy-behavioral-pattern-recognition-energy-and-comfort-management.

Dong, Bing, Zhaoxuan Li, Ahmad Taha, and Nikolaos Gatsis. 2018. "Occupancy-Based Buildings-to-Grid Integration Framework for Smart and Connected Communities." Applied Energy 219 (June): 123–37. https://doi.org/10.1016/j.apenergy.2018.03.007.

Dorokhova, Marina, Christophe Ballif, and Nicolas Wyrsch. 2020. "Rule-Based Scheduling of Air Conditioning Using Occupancy Forecasting." Energy and AI 2 (November): 100022. https://doi.org/10.1016/j.egyai.2020.100022.

Drgoňa, Ján, Javier Arroyo, Iago Cupeiro Figueroa, David Blum, Krzysztof Arendt, Donghun Kim, Enric Perarnau Ollé, et al. 2020. "All You Need to Know about Model Predictive Control for Buildings." Annual Reviews in Control 50 (January): 190–232. https://doi.org/10.1016/j.arcontrol.2020.09.001.

Dridi, Jawher, Manar Amayri, and Nizar Bouguila. 2022. "Transfer Learning for Estimating Occupancy and Recognizing Activities in Smart Buildings." Building and Environment 217 (June): 109057. https://doi.org/10.1016/j.buildenv.2022.109057.

Fang, Zhen, Nicolas Crimier, Lisa Scanu, Alphanie Midelet, Amr Alyafi, and Benoit Delinchant. 2021. "Multi-Zone Indoor Temperature Prediction with LSTM-Based Sequence to Sequence Model." Energy and Buildings 245 (August): 111053. https://doi.org/10.1016/j.enbuild.2021.111053.

Guan, Yepeng, and Yizhen Huang. 2015. "Multi-Pose Human Head Detection and Tracking Boosted by Efficient Human Head Validation Using Ellipse Detection." Engineering Applications of Artificial Intelligence 37 (January): 181–93. https://doi.org/10.1016/j.engappai.2014.08.004.

Happle, Gabriel, Jimeno A. Fonseca, and Arno Schlueter. 2020. "Context-Specific Urban Occupancy Modeling Using Location-Based Services Data." Building and Environment 175 (May): 106803. https://doi.org/10.1016/j.buildenv.2020.106803.

Hobson, Brodie W., H. Burak Gunay, Araz Ashouri, and Guy R. Newsham. 2020. "Clustering and Motif Identification for Occupancy-Centric Control of an Air Handling Unit." Energy and Buildings 223 (September): 110179. https://doi.org/10.1016/j.enbuild.2020.110179.

Hong, Tianzhen, Jared Langevin, and Kaiyu Sun. 2018. "Building Simulation: Ten Challenges." Building

Simulation 11 (5): 871–98. https://doi.org/10.1007/s12273-018-0444-x.

Ivanko, Dmytro, Åse Lekang Sørensen, and Natasa Nord. 2020. "Selecting the Model and Influencing Variables for DHW Heat Use Prediction in Hotels in Norway." Energy and Buildings 228 (December): 110441. https://doi.org/10.1016/j.enbuild.2020.110441.

Jang, Jihoon, Jinmog Han, and Seung-Bok Leigh. 2022. "Prediction of Heating Energy Consumption with Operation Pattern Variables for Non-Residential Buildings Using LSTM Networks." Energy and Buildings 255 (January): 111647. https://doi.org/10.1016/j.enbuild.2021.111647.

Jin, Yuan, Da Yan, Xuyuan Kang, Adrian Chong, Hongsan-- Sun, and Sicheng Zhan. 2021. "Forecasting Building Occupancy: A Temporal-Sequential Analysis and Machine Learning Integrated Approach." Energy and Buildings 252 (December): 111362. https://doi.org/10.1016/j.enbuild.2021.111362.

Khalil, Mohamad, Stephen McGough, Zoya Pourmirza, Mehdi Pazhoohesh, and Sara Walker. 2021. "Transfer Learning Approach for Occupancy Prediction in Smart Buildings." In 2021 12th International Renewable Engineering Conference (IREC), 1–6. Amman, Jordan: IEEE. https://doi.org/10.1109/IREC51415.2021.9427869.

Mirakhorli, Amin, and Bing Dong. 2016. "Occupancy Behavior Based Model Predictive Control for Building Indoor Climate—A Critical Review." Energy and Buildings 129 (October): 499–513. https://doi.org/10.1016/j.enbuild.2016.07.036.

Ouf, Mohamed M., William O'Brien, and Burak Gunay. 2019. "On Quantifying Building Performance Adaptability to Variable Occupancy." Building and Environment 155 (May): 257–67. https://doi.org/10.1016/j.buildenv.2019.03.048.

Parise, Alec, Miguel A. Manso-Callejo, Hung Cao, and Monica Wachowicz. 2021. "Prophet Model for Forecasting Occupancy Presence in Indoor Spaces Using Non-Intrusive Sensors." AGILE: GIScience Series 2 (June): 1–13. https://doi.org/10.5194/agile-giss-2-9-2021.

Park, Hansaem, Dong Yoon Park, Byeongjoon Noh, and Seongju Chang. 2022. "Stacking Deep Transfer Learning for Short-Term Cross Building Energy Prediction with Different Seasonality and Occupant Schedule." Building and Environment 218 (June): 109060. https://doi.org/10.1016/j.buildenv.2022.109060.

Pinto, Giuseppe, Riccardo Messina, Han Li, Tianzhen Hong, Marco Savino Piscitelli, and Alfonso Capozzoli. 2022. "Sharing Is Caring: An Extensive Analysis of Parameter-Based Transfer Learning for the Prediction of Building Thermal Dynamics." Energy and Buildings 276 (December): 112530. https://doi.org/10.1016/j.enbuild.2022.112530.

Qiao, Qingyao, Akilu Yunusa-Kaltungo, and Rodger Edwards. 2020. "Hybrid Method for Building Energy Consumption Prediction Based on Limited Data." In 2020 IEEE PES/IAS PowerAfrica, 1–5. https://doi.org/10.1109/PowerAfrica49420.2020.9219915.

SafeGraph Inc. 2023. "Weekly Patterns | SafeGraph Docs." 2023. https://docs.safegraph.com/docs/weekly-patterns.

Sharma, Himanshu, Saptarshi Bhattacharya, Soumya Kundu, and Veronica A. Adetola. 2022. "On the Impacts of Occupancy Sensing on Advanced Model Predictive Controls in Commercial Buildings." Building and Environment 222 (August): 109372. https://doi.org/10.1016/j.buildenv.2022.109372.

Shen, Weiming, Guy Newsham, and Burak Gunay. 2017. "Leveraging Existing Occupancy-Related Data for Optimal Control of Commercial Office Buildings: A Review." Advanced Engineering Informatics 33 (August): 230–42. https://doi.org/10.1016/j.aei.2016.12.008.

Stjelja, Davor, Juha Jokisalo, and Risto Kosonen. 2022. "Scalable Room Occupancy Prediction with Deep Transfer Learning Using Indoor Climate Sensor." Energies 15 (6): 2078. https://doi.org/10.3390/en15062078.

Taylor, Sean J., and Benjamin Letham. 2018. "Forecasting at Scale." Peer Jpreprints.

Tekler, Zeynep Duygu, and Adrian Chong. 2022. "Occupancy Prediction Using Deep Learning Approaches across Multiple Space Types: A Minimum Sensing Strategy." Building and Environment 226 (December): 109689. https://doi.org/10.1016/j.buildenv.2022.109689.

Warushavithana, Menuka, Saptashwa Mitra, Mazdak Arabi, Jay Breidt, Sangmi Lee Pallickara, and Shrideep Pallickara. 2021. "A Transfer Learning Scheme for Time Series Forecasting Using Facebook Prophet." In 2021 IEEE International Conference on Cluster Computing (CLUSTER), 809–10. Portland, OR, USA: IEEE. https://doi.org/10.1109/Cluster48925.2021.00076.

Weiss, Karl, Taghi M. Khoshgoftaar, and DingDing Wang. 2016. “Transfer Learning Techniques.” In Big Data Technologies and Applications, edited by Borko Furht and Flavio Villanustre, 53–99. Cham: Springer International Publishing. https://doi.org/10.1007/978-3-319-44550-2_3.

Wu, Wenbo, Bing Dong, Qi (Ryan) Wang, Meng Kong, Da Yan, Jingjing An, and Yapan Liu. 2020. “A Novel Mobility-Based Approach to Derive Urban-Scale Building Occupant Profiles and Analyze Impacts on Building Energy Consumption.” Applied Energy 278 (November): 115656. https://doi.org/10.1016/j.apenergy.2020.115656.

Yang, Tao, Arkasama Bandyopadhyay, Zheng O’Neill, Jin Wen, and Bing Dong. 2022. “From Occupants to Occupants: A Review of the Occupant Information Understanding for Building HVAC Occupant-Centric Control.” Building Simulation 15 (6): 913–32. https://doi.org/10.1007/s12273-021-0861-0.

Zhao, Yang, Peter Tu, and Ming-Ching Chang. 2019. “Occupancy Sensing and Activity Recognition with Cameras and Wireless Sensors.” In Proceedings of the 2nd Workshop on Data Acquisition To Analysis, 1–6. DATA’19. New York, NY, USA: Association for Computing Machinery. https://doi.org/10.1145/3359427.3361911.