

Towards the Detection of Building Occupancy with Synthetic Environmental Data

MANUEL WEBER, Munich University of Applied Sciences
CHRISTOPH DOBLANDER, Technical University of Munich
PETER MANDL, Munich University of Applied Sciences

Information about room-level occupancy is crucial to many building-related tasks, such as building automation or energy performance simulation. Current occupancy detection literature focuses on data-driven methods, but is mostly based on small case studies with few rooms. The necessity to collect room-specific data for each room of interest impedes applicability of machine learning, especially data-intensive deep learning approaches, in practice. To derive accurate predictions from less data, we suggest knowledge transfer from synthetic data. In this paper, we conduct an experiment with data from a CO₂ sensor in an office room, and additional synthetic data obtained from a simulation. Our contribution includes (a) a simulation method for CO₂ dynamics under randomized occupant behavior, (b) a proof of concept for knowledge transfer from simulated CO₂ data, and (c) an outline of future research implications. From our results, we can conclude that the transfer approach can effectively reduce the required amount of data for model training.

Additional Key Words and Phrases: Building occupancy, occupancy detection, transfer learning, carbon dioxide, synthetic environmental data, mass balance equation, deep learning, convolutional neural network

1 INTRODUCTION

Room-level occupancy in buildings provides crucial information to numerous use cases regarding, amongst others, building automation or building energy performance simulation. We differentiate between *occupancy estimation*, i.e. the determination of an exact count of occupants in a room, and *occupancy detection*, the binary discrimination between presence and absence. In both categories, a large number of models has been proposed over the last years. These include rule-based, stochastic and data-driven models. Recently, machine learning approaches have gained increasing interest. A recent study [2] shows that 56% of the models published between 2004 and 2019 are data-driven, 15% use neural network techniques. In addition to different model types, diverse sensing technologies can be found in literature. We classify sensing methods into *intrusive* and *non-intrusive*. Intrusive methods include additional sensors, such as light barriers or thermal cameras. Non-intrusive methods use existing data to predict occupancy states. *Environmental sensing* of carbon dioxide rates, temperature etc. is considered non-intrusive, as climate sensors are already widely installed for the purpose of automated climate control. Furthermore, climate sensors are more privacy-preserving than cameras, and, in contrast to transition-based measuring (e.g. with light barriers), they do not accumulate errors during the day. Chen et al. (2017) propose a first deep learning model for building occupancy estimation based on environmental factors, outperforming several state-of-the-art models [4]. Deep

learning allows the detection of occupancy at high precision without manual feature engineering. The challenge is that deep learning models require large quantities of room-specific labelled training data, which is not applicable in practice. Due to the high effort for collecting ground truth occupancy data for several days or weeks, models that provide good performance in a few test rooms cannot easily be applied to a variety of other rooms. We propose to use *transfer learning* to address this impediment. In our work, transfer learning is referred to in a transductive sense: Given equal learning tasks $T_s = T_t$ for two domains (rooms) R_s and R_t with data drawn from different distributions within the same feature space, and given labeled data from the source domain R_s and unlabeled or partially labeled data from the target domain R_t , transfer learning improves the training of T_t using R_s and T_s . Some publications have already reported benefits of a knowledge transfer between different rooms. Arief-Ang et al. (2017) extended a previously published seasonal decomposition model with a transfer approach, and used data recorded from a university office to effectively improve a model trained for a cinema hall [1]. In a later approach [10], transfer learning was successfully applied with a recurrent neural network (RNN), which is favorable in terms of generalizability. To the best of our knowledge, no deep transfer learning approach for environmental sensing-based detection of occupancy has been investigated to date. Moreover, synthetic data has not yet been considered for transfer, although physical simulation is common practice in the field of building technology. Tools such as EnergyPlus, TRNSYS and IDA ICE provide simulations for various environmental variables with high accuracy [5]. Regarding synthetic occupancy data, Chen et al. (2018) proposed an agent-based occupancy simulator [3] to generate stochastic occupancy data in office buildings.

2 APPROACH

We propose to leverage synthetic data generated from simulations to lower the need for costly real world data in context of occupancy detection. We conduct two simulations:

- (1.) Occupancy Simulation: For occupant presence and possibly additional, relevant actions, such as window opening actions
- (2.) Physical Simulation: For environmental factors, such as CO₂

A simulation of occupants, and possibly of their actions (1.), provides the ground truth data that also serves as an input to a simulation of the physical environment (2.). In this work, we focus on CO₂ dynamics. Also other environmental factors such as temperature or humidity may be considered. A physical simulation requires metadata about the room, such as its volume and infiltration rate. Occupant simulations need to take into account the type of room, e.g. 2-person office or lecture room. Otherwise, occupant behavior may not be comparable.

Authors' addresses: Manuel Weber, manuel.weber@hm.edu, Munich University of Applied Sciences, Lothstr. 64, Munich, Germany, 80335; Christoph Doblender, christoph.doblender@in.tum.de, Technical University of Munich, Munich, Germany; Peter Mandl, peter.mandl@hm.edu, Munich University of Applied Sciences, Lothstr. 64, Munich, Germany, 80335.

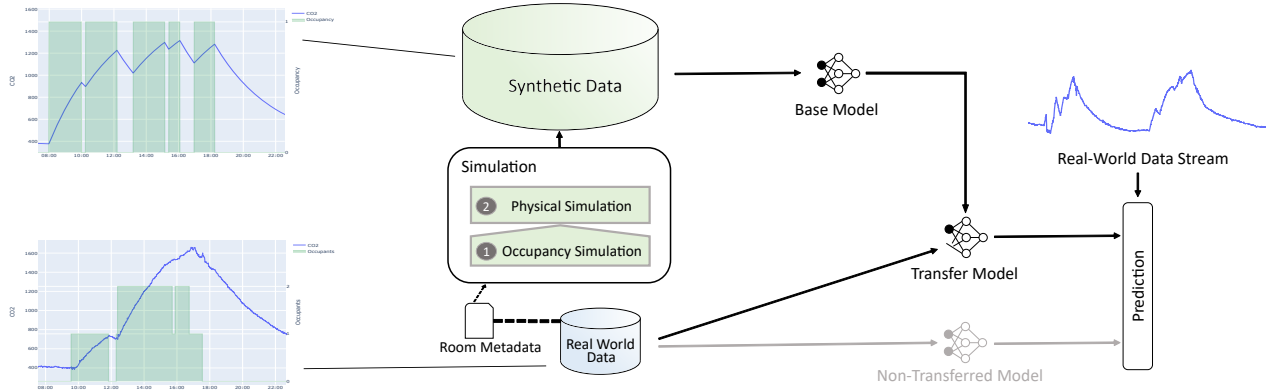


Fig. 1. Simulation-aided approach to occupancy detection

When building a model for a concrete room of interest, we distinguish two approaches:

- (a) run a preceding simulation based on metadata of the room of interest
- (b) use a general base model built from simulations on a variety of virtual rooms of the same room type

In this work, we use alternative (a) for a first proof of concept. It is illustrated in Fig. 1. The two consecutive simulations (occupancy and CO₂) generate a large-scale synthetic dataset. This dataset is used to pre-train a base model that is able to fit the general behavior of CO₂ dynamics under human presence or absence. Room-specific behavior due to certain conditions such as sensor placement, infiltration, window size etc. is then learned in a transfer step.

3 METHODS

This section describes the simulation methods we applied to simulate occupancy as well as CO₂ rates. In Subsec. 3.3 we then introduce the model architecture used to detect occupancy.

3.1 Occupancy Simulation

We simulate human occupancy in a naturally ventilated single office space of the size of a selected real world office. Occupancy behavior in offices can be subdivided into status transition events (e.g. arrival or final departure from work), random moving events (e.g. going to the bathroom) and meeting events [3]. As in [3], we combine the LIGHTSWITCH-2002 approach [8] for status transitions and Wang’s (2011) Markov chain approach [9] for random moving.

1. Status transition: For each day of simulation, we place an arrival at 8:00, a lunch break of 1h at 12:00, a departure at 18:00 and two 15-minute breaks at 10:00 respectively 15:00, with a random shift of ± 15 min for each event [8]. The timespans determine basic occupancy. Within these timespans, random moving events are simulated.

2. Random moving: Throughout all discrete time steps at basic occupancy, we use a Markov chain to successively determine the state of occupancy $occ[t] \in \{0, 1\}$ at time t , depending on the previous state $occ[t - 1]$ and a transition probability [9]. Transition

probabilities are picked from the transition matrix P depicted in Eq. 1. To consider a variety of different occupant types, we update P for each simulated day. The used sojourn times, s_0 for transitions from absence (0) to presence (1), and s_1 vice versa, are randomly selected within the following bounds.

$$P = \begin{bmatrix} 1 - (\frac{1}{s_0}) & \frac{1}{s_0} \\ \frac{1}{s_1} & 1 - (\frac{1}{s_1}) \end{bmatrix} \quad \begin{array}{l} 10min \leq s_0 \leq 60min \\ 30min \leq s_1 \leq 180min \end{array} \quad (1)$$

3. Window opening behavior: We decided to use a similar Markov chain to determine the window state $w[t] \in \{0, 1\}$ where 0 denotes that all windows or openings are closed, and 1 represents a state of ventilation. We set the time bounds to reflect different ventilation behavior to $60min \leq s_0^w \leq 8h$ and $5min \leq s_1^w \leq 30min$. Windows can be open also in case of absence due to random moving.

3.2 Carbon Dioxide Simulation

The change in CO₂ rate per time step within a room of volume V is calculated using Eq. 2. The formula is adopted from [6].

$$\frac{dc(t)}{dt} = \frac{\dot{m}(t)}{V} (c_{out} - c(t)) + \frac{G(t)}{V} \quad (2)$$

$c(t)$ is the indoor CO₂ concentration at time t , c_{out} is the outdoor CO₂ concentration, $\dot{m}(t)$ is the current mass flow rate, and $G(t)$ is the amount of CO₂ generated by human occupants present at the moment. Applying Eq. 2, we successively calculate CO₂ level values for time steps of one second, which can be later aggregated to the adequate granularity. For simplification purposes, air pressure and outdoor CO₂ are considered constant. The infiltration rate $\dot{m}(t)$ consists of a steady infiltration $\dot{m}_{inf}(t)$ independent from human occupancy, and a ventilation rate $\dot{m}_{vent}(t)$ determined by occupants’ actions, especially window opening events: $\dot{m}(t) = \dot{m}_{inf}(t) + \dot{m}_{vent}(t)$. To simplify, we use a constant value $\dot{m}_{inf}(t) = \dot{m}_{inf}$. Also, we consider a detailed calculation of $\dot{m}_{vent}(t)$ under consideration of varying window counts, -types and -sizes, as well as opening angles and air pressure differences, as unnecessary complex in a first proof of concept. Hence, whenever the window state is 1, we use a multiple

of \dot{m}_{inf} instead: $\dot{m}(t) = \dot{m}_{inf} \cdot vm$. A multiplier vm is randomly chosen from the range [10, 100] at each window opening event, and set to 1 if windows are closed. \dot{m}_{inf} is calculated from the volumetric air flow rate \dot{V}_{inf} by multiplication with the mass of air: $\dot{m}_{inf} = \dot{V}_{inf} \cdot m_{air}$. The value of m_{air} is set to 1.2754 g/l according to the IUPAC standard for the mass of dry air at standard temperature and pressure. \dot{V}_{inf} can be estimated by using CO₂ as a tracer gas, fitting a theoretical decay curve to the actual decay in CO₂ during a period of non-occupancy [6]. $G(t)$, defined in Eq. 3, is the product of the number of present occupants $n(t)$ (in our case 0 or 1), an average CO₂ generation rate of a single occupant g_{occ} , the mass of CO₂ m_{co2} and a unit translation term. For simulating intervals of one second, the term is divided by 60.

$$G(t) = n(t) \cdot g_{occ} \cdot m_{co2} \cdot \frac{1000}{60}, \quad n(t) \in \{0, 1\} \quad (3)$$

We choose $g_{occ} = 0.24$ l/min, which is the mean value measured by [7] for a person of an age between 21 and 28 standing at a desk. 0.18, the other value from [7], achieves lower results in our experiment.

3.3 Occupancy Detection

To detect occupancy, we apply the model architecture from [4], see Fig. 2. It combines a convolutional network with a deep bidirectional long short-term memory (DBLSTM). We downsample the data to a 1-min-granularity aggregating by the mean, and then use a sliding 15min-window on the input data stream. A one-dimensional convolutional layer (1D-Conv) and a max-pooling layer perform an automated feature extraction. As in [4], we use a filter size of 3, and a pooling factor of 2. We reduce the number of filters to 10, as we use CO₂ as a single input factor. The number of neurons in the following bidirectional long short-term memory (BLSTM) layers (200, 150 and 100 neurons) and the fully connected (FC) layers (300 and 200 neurons) are selected as proposed in [4], as well as the masking probabilities (0.5 and 0.3) for the dropout applied before the first and second FC layer for regularization purposes. A final softmax activation layer discriminates between presence (1) and absence (0).




Fig. 2. Deep learning model architecture based on [4]

4 EXPERIMENTAL SETUP

4.1 Datasets

Real world dataset: For a preliminary study, we measured the CO₂ rate in a 2-person office of a German university building located in Munich throughout 7 complete working days. A *Sensirion SCD30* non-dispersive infrared (NDIR) sensor was placed in the center of the room, and a measured CO₂ value was stored every 5 seconds.

According to the datasheet, the sensor provides a resolution of 1 ppm with an accuracy of ± 30 ppm +3%. The CO₂ data was collected between 10 and 19 March 2020. Occupancy ground truth data was manually recorded by the occupants over the observation period. **Simulated dataset:** In addition to this, we conducted a CO₂ time series simulation according to the methods previously described in Sec. 3.1 and 3.2 for a total of 500 artificial working days in a room with the same volume as the real world office ($V = 77.5$ m³). A natural infiltration rate and outdoor CO₂ concentration were estimated fitting Eq. 2 to observed CO₂ values from one unoccupied night not used for the dataset by minimizing the mean squared error (MSE). Consequently, they were set to realistic values of $\dot{V}_{inf} = 0.0046$ m³/s and $c_{out} = 360$ ppm. Table 1 summarizes the properties of the two datasets. Note that two occupants share the office, and guests may visit. Our simulation is only intended to provide basic information for a single-person scenario.

Table 1. Dataset Overview

	Real World Dataset	Simulated Dataset
Dataset Size	7 working days	500 working days
Time Granularity	5 sec	1 sec
Occupancy Values	[0, 3], mostly [0, 2]	[0, 1]
Presence Rate	29.28%	25.73%
CO ₂ Value Range	[338, 1749] ppm	[360, 1483] ppm

4.2 Experiments

We conducted a first study to show the positive effects of transfer learning from synthetic CO₂ data. First, 400 days of our simulated dataset were used to train a synthetic base model. The remaining 100 days were used to report the performance of the base model on simulated data. Using only small amounts of real world training data, we then trained two models: (1) a conventional deep learning model as described in Sec. 3.3, and (2) a transfer model with equal architecture and hyperparameters, using the same real world data, but based on the synthetic base model. The conventional model was trained using a uniform random initializer for model weights. Transfer learning was carried out by using the model weights of the base model for initialization instead, and retraining all layers. We evaluated with amounts of 1, 2, 3 and 4 days of training data. A cross validation was applied, using each of the 7 days in our dataset (and accordingly each 2, 3 or 4 consecutive days) for training in one iteration, and the remaining days as test data. Each iteration was repeated 10 times with a new seed value for randomization in initial weight generation and shuffling of input sequences. All experiments were carried out on an *Nvidia GeForce Tesla V100 SXM2* GPU, using the TensorFlow framework with Keras. For optimizer, learning rate, batch size and validation split, we chose RMSprop, 0.001, 70 and 0.2. Early stopping was applied after 20 epochs of no improvement in validation loss. To compare the results, we calculate the detection accuracy, which is defined as the number of correctly predicted occupancy states divided by the total number of predictions. As a second metric, we use the F1 score.

Table 2. Comparison of transfer model, non-transferred model and LR baseline for different training data extents*

	Training Data	1 Day	2 Days	3 Days	4 Days
Transfer Model	Accuracy	0.875 (± 0.057)	0.915 (± 0.020)	0.929 (± 0.014)	0.933 (± 0.015)
	F1 Score	0.764 (± 0.171)	0.863 (± 0.035)	0.888 (± 0.027)	0.898 (± 0.028)
	Epochs	27.2 (± 16.6)	26.1 (± 14.4)	22.6 (± 14.1)	19.9 (± 11.9)
Non-Transferred Model	Accuracy	0.715 (± 0.086)	0.874 (± 0.040)	0.895 (± 0.035)	0.914 (± 0.029)
	F1 Score	0.565 (± 0.162)	0.791 (± 0.065)	0.821 (± 0.071)	0.859 (± 0.053)
	Epochs	91.3 (± 54)	123 (± 50.2)	100.8 (± 46)	105.2 (± 34.4)
Logistic Regression (LR)	Accuracy	0.860 (± 0.073)	0.879 (± 0.058)	0.900 (± 0.035)	0.920 (± 0.016)
	F1 Score	0.767 (± 0.112)	0.803 (± 0.080)	0.829 (± 0.056)	0.866 (± 0.038)

*mean values (\pm standard deviations), best scores are highlighted in bold

5 RESULTS

Table 2 reports the mean and standard deviation values of accuracies, F1 scores and training epochs under different extents of training data, in both the transfer and non-transfer setting. Additionally, results are compared to a logistic regression (LR) baseline. The pre-trained transfer model clearly outperformed the non-transferred model, with a higher impact the less training data was used. Vice versa, only half of the training data was required to achieve a similar accuracy as in the non-transfer setting: 0.87 was reached with one instead of two days, and 0.91 with two instead of four days of training data. Besides a substantial improve in accuracy and F1 score, also the standard deviation was reduced in most cases. Hence, the transfer approach can also improve model robustness. As data was scarce, the deep learning model was unable to reach its full potential, and performed even slightly worse than an LR classifier. The transfer model, in contrast, showed superior results in nearly all cases. Regarding training times, the transfer model was able to be trained in only a fraction of the training epochs. The number of epochs reported in Table 2 indicate after how many training epochs validation loss reached a minimum with respect to all previous and 20 subsequent epochs. By average, 91 epochs were trained for the non-transferred model with one day of training data, and above 100 with multiple days of training data. In contrast to this, transfer learning allowed a validation loss convergence after less than 30 epochs. However, this reduction of training times results from the additional upfront effort for pre-training. The effect is only beneficial if the base model is reused for multiple rooms. Pre-training required 30 epochs on data from 400 synthetic training days. The resulting base model achieved an accuracy of 0.981 and an F1 score of 0.963 on 100 simulated days. Without the subsequent transfer step, the base model by itself was not able to make accurate predictions on the real world test dataset. Simulations do not replace but reduce the need for data collection.

6 CONCLUSION & FUTURE RESEARCH

In our experiment, we have demonstrated that transfer from synthetic data can effectively improve model performance and robustness regarding occupancy detection. The results encourage the use of simulations in this field of limited real world data to enable deep learning in practice. We see future research in finding a concrete

method considering how to generate adequate synthetic data, and how to accomplish transfer. It should also be investigated whether a large, generalized base model, as it is common in the field of image processing for instance, may be beneficial for a variety of different rooms. We aim to prepare a broader base model and evaluate on multiple room types. Furthermore, we are planning to investigate other model architectures and their ability to transfer knowledge. For the purpose of demonstration, in this work several constant values and arbitrary boundaries were used in the simulation. We want to overcome these limitations, and also consider predicting the number of present occupants.

ACKNOWLEDGMENTS

Special thanks go to the *OpenPower@TUM* project and the industry partner IBM for providing us with the computing infrastructure for our experiments.

REFERENCES

- [1] I. B. Arief-Ang, F. D. Salim, and M. Hamilton. 2017. DA-HOC: Semi-Supervised Domain Adaptation for Room Occupancy Prediction using CO2 Sensor Data. *Proceedings of the 4th ACM International Conference on Systems for Energy-Efficient Built Environments (BuildSys)* (2017), 1–10.
- [2] S. Carlucci, M. de Simone, S. K. Firth, M. B. Kjærsgaard, R. Markovic, M. S. Rahaman, M. K. Annaqeeb, S. Biandrate, A. Das, J. W. Dziedzic, G. Fajilla, M. Favero, M. Ferrando, J. Hahn, M. Han, Y. Peng, F. Salim, A. Schlüter, and C. van Treeck. 2020. Modeling occupant behavior in buildings. *Building and Environment* 174 (2020), 106768.
- [3] Y. Chen, T. Hong, and X. Luo. 2018. An agent-based stochastic Occupancy Simulator. *Building Simulation* 11, 1 (2018), 37–49.
- [4] Z. Chen, R. Zhao, Q. Zhu, M. K. Masood, Y. C. Soh, and K. Mao. 2017. Building Occupancy Estimation with Environmental Sensors via CDBLSTM. *IEEE Transactions on Industrial Electronics* 64, 12 (2017), 9549–9559.
- [5] D. Mazzeo, P. Romagnoni, N. Matera, G. Oliveti, C. Cornaro, and L. De Santoli. 2019. Accuracy Of The Most Popular Building Performance Simulation Tools: Experimental Comparison For A Conventional And A PCM-Based Test Box. *Proceedings of the 16th IBPSA Conference* (2019), 4530–4537.
- [6] P. Parsons. 2014. Determining Infiltration Rates and Predicting Building Occupancy Using CO2 Concentration Curves. *Journal of Energy* (2014), 1–6.
- [7] C. Reiff, K. Marlatt, and D. R. Dengel. 2012. Difference in caloric expenditure in sitting versus standing desks. *Journal of physical activity & health* 9, 7 (2012), 1009–1011.
- [8] C. F. Reinhart. 2004. Lightswitch-2002: a model for manual and automated control of electric lighting and blinds. *Solar Energy* 77, 1 (2004), 15–28.
- [9] C. Wang, D. Yan, and Y. Jiang. 2011. A novel approach for building occupancy simulation. *Building Simulation* 4, 2 (2011), 149–167.
- [10] T. Zhang and O. Ardakanian. 2019. A domain adaptation technique for fine-grained occupancy estimation in commercial buildings. *Proceedings of the 2019 Internet of Things Design and Implementation (IoTDI)* (2019), 148–159.