

Enabling Building Energy Auditing Using Adapted Occupancy Models

Ankur Kamthe, Varick Erickson, Miguel Á. Carreira-Perpiñán and Alberto Cerpa

Electrical Engineering and Computer Science - University of California, Merced
{akamthe,verickson,mcarreira-perpinan,acerpa}@ucmerced.edu

Abstract

Understanding building energy consumption has become important due to stricter energy regulations, increasing energy costs and also as buildings have long term impact on energy consumption. In order to recommend retrofits, it is important to have accurate estimates for building energy consumption which is affected significantly by occupancy patterns. This paper explores the development of static occupancy models using a model adaptation technique that is able to capture accurately features of occupancy distributions typically found using a large amount of training data (days, weeks, months). Using only one day of training data that can be easily recorded without any infrastructure but battery-operated sensors with on-board memory, we show that our adapted occupancy model can estimate energy savings of 10.9%; and the room temperatures for the adapted model schedules were 0.5°F and 1.4°F off from the target temperatures for summer and winter months, respectively. This performance was on par with models trained with four times as much data. Our proposed technique can be used by energy auditors to estimate energy savings for existing buildings and by building energy managers to optimize static schedules which assume maximum occupancy.

Categories and Subject Descriptors

I.6.5 [Simulation and Modeling]: Model Development

General Terms

Algorithms, Machine Learning, Measurement

Keywords

Model Adaptation, Occupancy, HVAC, Energy Saving

1 Introduction

In 2010, residential and commercial buildings accounted for 40.4% of the total U.S. primary energy consumption [1].

Also, the heating, cooling and ventilation systems were responsible for approximately 50% of the total building energy usage [1]. In addition, a vast majority of existing buildings are older than 20 years and do not meet current energy efficiency construction standards [12]. So, not only do buildings account for a significant portion of energy consumption, but also they impact the long-term energy consumption. Therefore, a significant reduction in building energy consumption will directly impact the overall energy usage numbers.

A common way to estimate energy savings is through building energy auditing. A detailed energy audit [12] is comprehensive, time-consuming and involves inspection and analysis of the energy consumption from utility bills. In addition, sensors are deployed on-site to measure and verify the energy use (lighting, office equipment, fans, chillers, etc.). Typically, for small buildings on-site work takes one day, while bigger buildings might require two days. This data is used by computer simulation programs such as DOE-2, EnergyPlus, etc. to evaluate and recommend energy retrofits.

Regulating ventilation based on occupancy data has been shown to reduce HVAC energy usage by 10-15% [4]. Studies have optimized HVAC system usage based on occupancy [5, 15, 2]. In all these approaches, the core idea is to instrument the building extensively to infer occupancy data. Other studies [7, 15, 8] have created models from the occupancy data and have used them to estimate energy savings. The drawback is the need for large training datasets (weeks, months) to compute reliable estimates for the numerous model parameters. Also, once a model is created, it is not possible to directly port it to a different building due to different floor-plans, room functions and the resulting occupancy patterns. Therefore, it implies that for modeling new buildings, it is required to record occupancy traces for extended periods of time (weeks, months) using extensive sensor network infrastructure comprising of sensor nodes, data logging units, power cables, etc. This prevents applications such as building energy auditing from making use of occupancy data to maximize the energy savings of a new building using short occupancy data traces (1-2 days).

For such situations, it makes sense to use a reference building occupancy model that has been trained with extensive data and adapt it to the new building given a far smaller occupancy data trace than would be necessary to train a new model from scratch. Such occupancy datasets can be recorded using few sensors with the data being stored in the

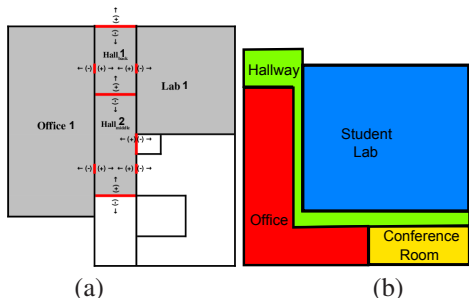


Figure 1. Building floorplans for reference model dataset (a) and adaptation model dataset (b), respectively.

on-board memory. Also, there is no need to install cabling to power these nodes as they can survive for a single day using their battery packs. This will enable highly accurate buildings energy auditing based on usage patterns.

In this paper, we try to address the problem posed by the large data requirement for training occupancy models. We propose a technique to adapt the parameters from a reference building occupancy model to fit the smaller occupancy dataset for the new building.

2 Background and Motivation

2.1 Occupancy Modeling Approaches

Building energy simulation tools such as eQuest and EnergyPlus use predefined occupancy profiles for office buildings based on day-type and maximum room occupancy. Recent studies have optimized energy consumption by adjusting the HVAC based on occupancy data. Several approaches [6, 15, 16, 2] have proposed aggregating data from multiple sensor streams to infer occupancy data. These approaches involve measurement of electrical loads, extensive deployment of sensors such as CO₂ and passive infra-red (PIR) sensors and computer network activity monitoring for determining whether a room is occupied. Other approaches have utilized PIR and door sensors, which give a binary indication of occupancy to optimally ventilate rooms in office buildings [2]. In [7, 8], occupancy models were proposed which were estimated using data from a camera sensor network. In the aforementioned approaches, the studies assumed/required the presence of IP-based connectivity for a back-channel to transmit and store data, electronic door locks which could provide access logs, submetering to monitor the consumption within a building and dense camera sensor network deployment. Also, to enable their models to predict occupancy in advance to optimally control lighting and ventilation, the data streams need to be recorded over days, weeks and months to incorporate occupancy data to improve building conditioning by doing occupancy prediction. For example: 5 day dataset using a camera sensor network [7, 8], 3 month dataset from contact and PIR sensors, network activity, building energy usage [15] and 10 month dataset from multiple streams (room maps, appliance power meters, building sub-metering) [16].

2.2 Model Adaptation Techniques

Similar to occupancy modeling, in the area of speech recognition, training word recognition systems involve estimating parameters for Gaussian-mixture based HMMs,

which have a very high number of states and mixture components. For example: in [13], the HMM had 1778 states with 6 mixture components per state. The training process to estimate good parameter estimates for such models required huge amounts of training data. Model adaptation techniques such as Maximum A-Posteriori (MAP) [9] and Maximum Likelihood Linear Regression (MLLR) [13] have been proposed to solve the training data problem. The drawback of MAP is that it can only update the parameters seen in the training data. MLLR is based on tying together groups of parameters (means, covariances) and using linear transformations of the parameters. It significantly outperforms MAP with very little data. In [10], the means of a mixture of multivariate Bernoulli (MMB) distribution were tied using a non-linear transformation for modeling the MNIST handwritten digits dataset and 802.15.4 wireless traces. Thus, parameter-sharing based techniques have proven their utility for solving a wide variety of modeling problems that required large training datasets.

In this paper, we apply these well-studied techniques to solve the large data requirement problem of occupancy modeling. Our approach for estimating occupancy models makes the occupancy data collection process practical and easy to carry out for practitioners such as building energy auditors.

2.3 Motivation

Figure 1 shows parts of two different building floorplans. The occupancy patterns for areas within these floorplans is shown in Figure 2. As the occupancy patterns are different over different floorplans, it is not possible to port occupancy models created for one floorplan to a completely different floorplan. Then the only way to create a new model is to deploy a sensing infrastructure to record the occupancy patterns. However, most of the current occupancy estimation approaches are ill-suited for existing buildings that lack infrastructure such as submetering systems and electronic locking systems among others. In such situations, the importance of wireless sensor networks (WSNs) is realized as they allow for sensing without any pre-existing infrastructure. However, the limitations of a WSN, namely limited storage and limited battery life, are exposed when deployed to record occupancy data. This is because when there is no network access, the sensor nodes cannot transmit the data and hence, need to store it on-board. PIR sensor have long lifetimes but give binary indication of occupancy, whereas cameras sensors are capable of accurate occupancy estimates [17, 11] but quickly deplete the limited battery-life of sensor nodes. This imposes a practical limit on the amount of data that can be recorded using a sensor network to train occupancy models. In the context of energy auditing where data collection is limited to 1-2 days [12], there is need for models that can provide reliable, or at a minimum, conservative estimates of occupancy using 1 or 2 days of training data as input.

3 Proposed Approach

In general, model adaptation techniques adapt the parameters of a pre-existing model, called reference model, to match the distribution of the new data (adaptation data) from a small dataset. In this paper, we present a technique to adapt the parameters of the multivariate Gaussian model

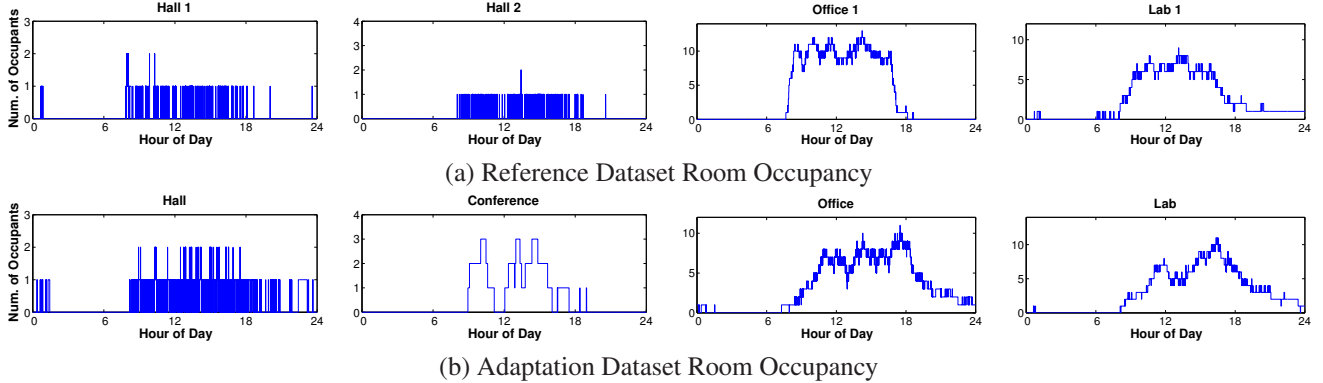


Figure 2. Room occupancy averaged over the length of dataset (5-days) for every hour for the reference model (a) and adaptation (and retrained) model (b), respectively.

(MVGM) [7]. Our reference model is an MVGM trained using an extensive 5-day occupancy dataset from building with floorplan as shown in Figure 1(a). The occupancy data that we record for the new building floorplan (see Figure 1(b)) is called adaptation data. The variation in room occupancies for the reference model dataset and adaptation dataset is shown in Figure 2. The reference model dataset rooms show occupancy patterns that are drastically different from the adaptation dataset rooms. The room occupancies for both datasets were recorded using a camera sensor network [11].

In this section, we briefly describe the MVGM, explain our proposed approach for adapting the MVGM parameters using a non-linear transformation and provide a generalized EM algorithm to estimate the transformation parameters.

3.1 Multivariate Gaussian Model (MVGM)

In [7], Erickson et al. have proposed a multivariate Gaussian model (MVGM) for modeling occupancy data. A separate multivariate Gaussian distribution was trained for every hour of the day using occupancy data from several rooms of a building floorplan. A natural extension of the MVGM is to use a mixture of multivariate Gaussians in place of a single multivariate Gaussian. This gives the model additional flexibility to model the room occupancy patterns for every day of the week for each hour. In this paper, we use a mixture of multivariate Gaussian distribution for modeling occupancy for each hour. Let $X_h = (x_{h1}, \dots, x_{hN})$ denote a series of N ($=3600$, one per second) occupancy data observations collected during hour h of a day. Each x_{hn} is a vector of dimension D , where D is the number of rooms. The probability density is given by:

$$p(\mathbf{x}_{hm}) = \sum_{m=1}^M \pi_{hm} p(\mathbf{x}_{hm}|m)$$

and

$$p(\mathbf{x}_{hm}|m) = \frac{1}{(2\pi)^{\frac{D}{2}} |\Sigma_h|^{-\frac{1}{2}}} \exp \left\{ -\frac{1}{2} (\mathbf{x}_{hm} - \mu_{hm})' \Sigma_h^{-1} (\mathbf{x}_{hm} - \mu_{hm}) \right\}$$

where there are M components and the parameters are the mixing proportions π_{hm} (which are positive and sum to one), the means μ_{hm} ($\mu_{hm} \in \mathbb{R}^D$, $\mu_{hm} > 0$) and the covariance matrix Σ_h . From this, the average occupancy vector for any given hour can be computed as follows:

$$\mu_h = \sum_{m=1}^M \pi_{hm} \mu_{hm}$$

The parameter estimation for a mixture of multivariate Gaussian distribution is done with an EM algorithm [3]. EM is an iterative algorithm that alternates between an expectation (E) step and a maximization (M) step, till the algorithm converges to a local minima. In the E step, the posterior probability of each component given a data vector is computed using the current parameter. The M step estimates updated parameters for each component using the posterior probabilities. In the context of adaptation, we will call *re-training* the process of estimating the mixture of multivariate Gaussian parameters using this EM algorithm given the adaptation data, and initializing the parameters to those of the reference model. Retraining with little data results in over-fitting. Such models generalize poorly to future data. In case of occupancy modeling, when little adaptation data is available, some rooms may be unoccupied on the day the data was recorded. The corresponding μ_{hm} will clamp to 0, indicating that the room remains unoccupied at all times, whereas the room may be heavily utilized on other days.

4 Adapting the MVGM

Suppose we have a MVGM (the *reference* model) corresponding to a building floorplan for every hour ($h \in \{1, \dots, 24\}$) of the day. The reference model for each hour is trained on an extensive occupancy dataset resulting in good parameter estimates for the mean (μ_m), covariance Σ (same for all components) and mixture proportions π_m where $m \in \{1, \dots, M\}$ (h is dropped from subscript to simplify notation). We are now given an adaptation occupancy dataset comprised of N D -dimensional vectors for a different building layout for which we need to estimate the adapted MVGM parameters, $\{\tilde{\pi}_m, \tilde{\mu}_m\}_{m=1}^M$ and $\tilde{\Sigma}$. Our adaptation algorithm is based on the idea of tying the MVGM parameters, specifically the means, through a transformation of the reference MVGM means. Adapting with a linear transformation can restrict the changes in some of the $\tilde{\mu}_{md}$, because when the smallest or largest room occupancy value reaches saturation, it prevents less extreme occupancy values from adapting significantly (since all rooms share the same transformation within a component of the Gaussian mixture). Adapting with a sigmoid transformation avoids this problem. This strategy is similar to the approach in [10], where Bernoulli parameters must lie in $[0, 1]$. In our proposed approach, we




Figure 3. Occupancy models for Office, Lab and Conference rooms using retrained MVGM trained with 4 days of data (MVGM-R4), adapted MVGM trained with 1 day of data (MVGM-A1), and the corresponding retrained model MVGM-R1.

tie the means of the reference model together using a scaled-sigmoid transformation as follows; $\tilde{\mu}_{md} = \sigma(\mu_{md}; a_m, b_m) = \frac{O_m}{1 + e^{-(a_m \mu_{md} / O_m + b_m)}}$, $d = 1, \dots, D$

where, O_m is the maximum room occupancy out of all rooms and is known a priori. The transformation has only two parameters a_m, b_m per component, which enables us to estimate good values from the small adaptation occupancy dataset. As a_m, b_m are shared among the D -reference model means, their effect is propagated among all the component means. The mixture proportions are considered free during adaptation, subject to constraint that they sum to 1. The covariance for the mixture distribution is kept fixed ($\tilde{\Sigma} = \Sigma$) during the adaptation. Thus, our algorithm needs to maximize the likelihood of the adaptation data over $3M - 1$ parameters ($\tilde{\pi}_1, \dots, \tilde{\pi}_{M-1}$ and $a_1, b_1, \dots, a_M, b_M$), which are small in comparison to the $(D + 1)M + D^2 - 1$ parameters (means, covariance matrix and mixing proportions) for the retrained MVGM.

We solve for a_m, b_m for the adapted MVGMs for each hour using a generalized EM algorithm. The objective function is the log-likelihood of the adaptation data given the constrained MVGM with $3M - 1$ free parameters:

$$L(\{\tilde{\pi}_m, a_m, b_m\}_{m=1}^M) = \sum_{n=1}^N \log \sum_{m=1}^M \tilde{\pi}_m p(\mathbf{x}_n; a_m, b_m)$$

where $p(\mathbf{x}_n; a_m, b_m)$ is a multivariate Gaussian where $\tilde{\mu}_{md} = \sigma(\mu_{md}; a_m, b_m)$. We maximize the objective function using a generalized EM (GEM) algorithm [14], which resembles the one in [10]. We initialize from the reference model, setting $\tilde{\pi}_m = \pi_m$, and use an identity transformation (approximated using a sigmoid with $a_m = 6, b_m = -3$) to set $\tilde{\mu}_{md} = \mu_{md}$. The GEM algorithm stops iterating between the E and M steps after converging to a local minima (estimates of a_m, b_m, π_m where $(m \in \{1, \dots, M\})$).

5 Modeling Results

To evaluate the performance of our proposed approach, we compare the adapted models with the retrained models in terms of model log-likelihood. Out of the 5-day adaptation dataset, we used 4 days for training using the retraining and adaptation, and the 5th day was used as the test set.

Figure 4 shows the variation in test set log-likelihood of the retrained and the adapted MVGMs as a function of the amount of adaptation data. With just 1 day of data, the log-likelihood of the adapted model (MVGM-A1) is already significantly better than the corresponding retrained model (MVGM-R1). As we increase the adaptation data, the log-likelihood of the adapted model approaches the optimal log-likelihood (MVGM-R4). The retrained model needs

much more data to achieve a comparable log-likelihood to the adapted model; and its performance is much more variable depending on the specific adaptation vectors used (large error bars). Figure 3 displays the occupancy models for every hour for each room. We can observe that for rooms such as the Office and Lab that are heavily utilized, there is very little difference between the 3 models. However, for the conference room (Figure 3), we can that MVGM-A1 is able to capture the occupancy behavior similar to MVGM-R4 and also the average occupancy as shown in Figure 2. This is in stark contrast to MVGM-R1 that overfits to the adaptation data used for training the model, where the conference room was occupied only between 18:00-20:00 hours.

This clearly demonstrates that the adapted model generalizes to data in the unseen test-sets much better than the retrained model with significantly less training data.

6 Building Energy Simulation Results

Using the MVGM-R4, MVGM-R1 and MVGM-A1 models, we construct an expected occupancy schedule for each hour using the average occupancy vector. For simulating the building energy consumption, we use these occupancy schedules to estimate ventilation and temperature schedules. We inputted these schedules into an EnergyPlus model of the building floorplan (total 32,000 sq.ft.) from which we have adaptation data for a Hall, Office, Lab and Conference room (approx. 12,000 sq.ft., see Figure 1). The building is complaint to applicable codes and has a single duct terminal-reheat HVAC system. To estimate the energy savings and conditioning effectiveness, we compared to a baseline strategy. The baseline strategy conditions rooms assuming maximum room occupancy between 7 a.m. and 10 p.m. and is off at other times. Also, we compare with OBSERVE [8]. It utilizes a Markov chain, in conjunction with a camera sensor network, to model the temporal changes in occupancy of a building. In OBSERVE, the camera sensor network sources the data to contrast and correct the transition probabilities of the Markov chain. OBSERVE has shown significant energy savings and close to optimal conditioning.

6.1 Energy Savings

Figure 5 shows the energy savings. Note, only the Hall, Office, Lab and Conference rooms (12,000 sq.ft.), for which adaptation data was available, were conditioned using MVGM-R4, MVGM-R1, MVGM-A1 and OBSERVE. The other parts of the floorplan (20,000 sq.ft.) were conditioned using the baseline strategy. In our EnergyPlus simulations, energy consumption was computed for the major HVAC components; the fans, heating (gas), and cooling




Figure 4. Log-likelihood of the different models as a function of the days in the adaptation dataset.




Figure 5. Estimated energy savings using different strategies. OBSERVE and MVGM-R4 perform similarly. MVGM-A1 is conservative whereas MVGM-R1 overestimates savings

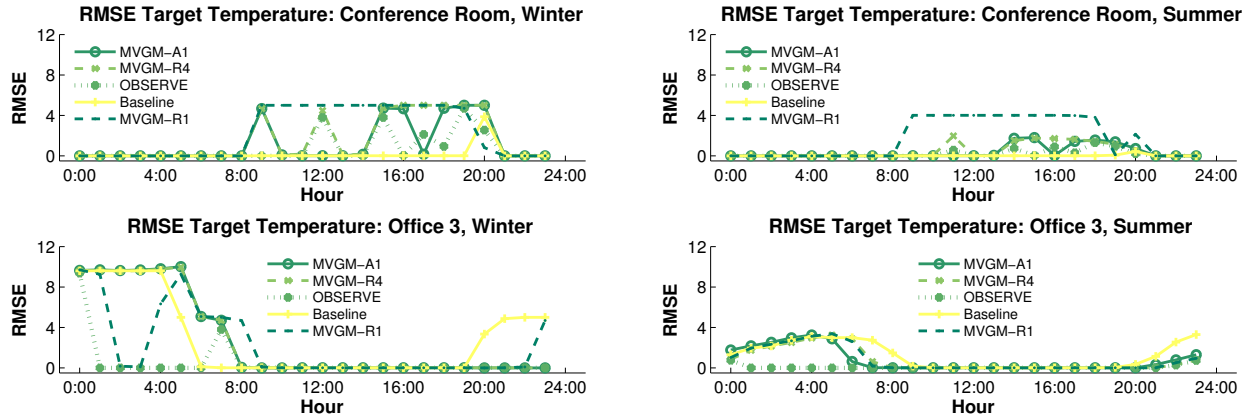


Figure 6. Comparison of the models with respect to temperature in the Conference room and Office 3 for the summer and winter seasons. The root mean square error of the target and measured temperature is examined for each hour.

(gas). OBSERVE and MVGM-R4 which are trained using 4 days of occupancy data show average annual energy savings of 11.2% and 11.4%, respectively. In comparison, our proposed approach MVGM-A1, which used only 1 day of occupancy data, shows annual savings of 10.9%. MVGM-R1 estimates annual savings of 12.9%. This is explained by the fact that the MVGM-R1 model parameters (Figure 3) indicate that the conference room is mostly empty compared to MVGM-R4 and MVGM-A1. OBSERVE predicts occupancy using its Markov chain model and also corrects the ventilation and temperature using occupancy data updates from the sensor network. Therefore, the energy savings from OBSERVE are the most that one can expect to achieve. Our proposed approach, MVGM-A1 is conservative in its energy saving estimate, whereas MVGM-R1 over-estimates savings due to overfitting of the model parameters.

6.2 Conditioning Effectiveness

There are two different conditioning criteria that need to be examined: thermal comfort (ASHRAE Std. 55) and outside air ventilation effectiveness (ASHRAE Std. 62.1). We compare the different models on the test day occupancy conditions for the winter and summer seasons.

6.2.1 Thermal Comfort

The optimal temperatures for the heating and cooling schedules used in the analysis are 78°F and 75°F respectively. Figure 6 shows the root mean squared error (RMSE) of the room temperature to the target temperature during occupied periods for the Conference Room and Office for the winter and summer seasons. The Baseline strategy is optimal between 7:00-22:00 hours as it conditions rooms between

those times and is off at other times. OBSERVE has the best overall performance of the four models with lower RMSE for almost all hours for both the winter and summer. This is because OBSERVE has perfect visibility of the room occupancy based on data from the camera sensor network. In contrast, the models (MVGM-R4, MVGM-R1 and MVGM-A1) have static temperature schedules created using the mean occupancy of each model at every hour.

For Office and Lab, all models do equally well. This is because, the adaptation data used for training the models is not significantly different from the test day. However, for the conference room, which is mostly unoccupied, except for certain hours on certain days, we can see the deviation from the Baseline. In summer, the RMSE of the Baseline, MVGM-R4, OBSERVE and MVGM-A1 models is below 0.5°F whereas the RSME for MVGM-R1 is 1.8°F. In winter, the RMSE of the Baseline, MVGM-R4, OBSERVE and MVGM-A1 models is below 1.4°F whereas the RSME for MVGM-R1 is 2.4°F. The MVGM-R1 model conditions the conference room only starting from 18:00 hours till 20:00 hours, as the corresponding schedule indicates occupancy between those hours. This explains the worse behavior of MVGM-R1. The RMSE error is higher for the winter than the summer because the heating plant used for the simulation is less effective than the cooling plant.

6.2.2 Ventilation

Figure 7 shows a comparison of the ventilation rates. We computed the ventilation required for the area occupied by Hall, Lab, Office and Conference rooms. All ventilation strategies perform better than Baseline (maximum ventila-




Figure 7. Comparison of the ventilation rates. The ventilation rate is the sum of Hall, Lab, Office and the Conference Room ventilation rates.

tion from 7am-10pm); the Baseline rate exceeds the required by at least 23%. OBSERVE has perfect visibility and therefore, ventilates the best, exceeding required by 2.8%. The static models over and under ventilate at times. MVGM-R4, MVGM-R1 and MVGM-A1 all perform similarly, exceeding required by average of 8.3%. At times, when the models under-ventilate, it is by an average of 12 CFM (2 people) over a floorplan area of roughly 12,000 sq.ft.

7 Discussion

Our adapted model, MVGM-A1 outperforms the corresponding retrained model, MVGM-R1, trained with one day of adaptation data in terms on test-set log-likelihood. MVGM-A1 gives a conservative estimate of energy savings when compared to MVGM-R1, which over-estimates the energy savings. In terms on conditioning effectiveness, MVGM-A1 matches MVGM-R4 in being closer to target temperature compared to MVGM-R1. This is because, MVGM-R1 overfits to the adaptation data and estimates that the conference will only be occupied only between 18:00-20:00 hours for all days of the week, resulting in greater energy savings than all other models. In contrast, MVGM-A1 generalizes beyond the adaptation data, estimating that the conference room will be occupied more times than is observed in the limited adaptation data. This behavior can be attributed to the use of the parameter-sharing based model adaptation approach. As the reference model parameters are tied together with the transformation, the resulting adapted model parameters cannot clamp to 0 when the room is unoccupied. This results in the adapted model parameters having non-zero mean occupancy for rooms that are unoccupied in the adaptation data, which is closer to reality.

Typically, model adaptation is effective if the reference model is close enough to the new adaptation occupancy data, such that the transformation parameters can then map to the distribution of the adaptation data. As an analogy, for a English word-recognition system, the reference model from one English language speaker can be adapted to a different (new) English speaker with just a few utterances. However, this does not work when using a English speaker reference model for French word recognition system, i.e. adaptation data is from a French speaker. Therefore, in our paper, our reference and adaptation data were from two different floorplans but both were office buildings. Our results may be different if the adaptation data were from a building used exclusively for classrooms. For the purpose of energy auditing, it would be useful to have reference models for a few dif-

ferent types of buildings with fundamentally different usage patterns such as government offices, private offices, schools, research buildings, etc.

8 Summary

In this paper, we addressed the issue of estimating occupancy models when it is possible to collect only 1 or 2 days of occupancy data. Using model adaptation, we reduced the amount of data required for estimating good parameter values for the occupancy models to 1 day. The adapted model generalizes to occupancy conditions not recorded in the small training dataset, conservatively predicting 10.9% energy savings and conditioning effectiveness on par with other models that require 4 times as much training data.

Our model adaptation approach can enable highly accurate building energy auditing, by allowing auditors to estimate energy savings using the 1-2 days of occupancy data recorded on-site. Also, it is a useful tool for building energy managers to optimize their baseline strategies. They can collect occupancy data using a handful of wireless sensors over 1 day and use model adaptation to get reliable occupancy model estimates. In both these cases, apart from the battery-operated wireless camera sensor nodes, there is no need for additional infrastructure to infer occupancy data that can provide significant energy savings in return.

9 References

- [1] *2010 Building Energy Data Book*. U.S. Dept. of Energy, 2011.
- [2] Y. Agarwal, B. Balaji, S. Dutta, R. K. Gupta, and T. Weng. Duty-cycling buildings aggressively: The next frontier in HVAC control. In *IPSN'11*.
- [3] C. M. Bishop. *Pattern Recognition and Machine Learning*. Springer-Verlag, 2006.
- [4] M. J. Brandemuehl and J. E. Braun. The impact of demand-controlled and economizer ventilation strategies on energy use in buildings. *ASHRAE Transactions*, 1999.
- [5] D. T. Delaney, G. M. P. O'Hare, and A. G. Ruzzelli. Evaluation of energy-efficiency in lighting systems using sensor networks. In *BuildSys '09*.
- [6] R. Dodier, G. Henze, D. Tiller, and X. Guo. Building occupancy detection through sensor belief networks. *Energy & Buildings*, 2006.
- [7] V. Erickson, Y. Lin, A. Kamthe, R. Brahme, A. Cerpa, M. Sohn, and S. Narayanan. Energy efficient building environment control strategies using real-time occupancy measurements. In *BuildSys*, 2009.
- [8] V. L. Erickson, M. A. Carreira-Perpiñán, and A. E. Cerpa. OBSERVE: Occupancy-based system for efficient reduction of HVAC energy. In *IPSN'11*.
- [9] J.-L. Gauvain and C.-H. Lee. Bayesian learning for hidden Markov model with Gaussian mixture state observation densities. *Speech Communication*, 11(2-3):205-213, June 1992.
- [10] A. Kamthe, M. Á. Carreira-Perpiñán, and A. Cerpa. Adaptation of a mixture of multivariate bernoulli distributions. In *IJCAI 2011*.
- [11] A. Kamthe, L. Jiang, M. Dudys, and A. Cerpa. SCOPES: Smart cameras object position estimation system. In *EWSN*, 2009.
- [12] M. Krarti. *Energy audit of building systems: An engineering approach*. CRC Press, 2011.
- [13] C. J. Leggetter and P. C. Woodland. Maximum likelihood linear regression for speaker adaptation of continuous density hidden Markov models. *Computer Speech and Language*, 9(2):171-185, Apr. 1995.
- [14] G. J. McLachlan and T. Krishnan. *The EM Algorithm and Extensions*. John Wiley & Sons, second edition, 2008.
- [15] G. R. Newsham and B. J. Birt. Building-level occupancy data to improve ARIMA-based electricity use forecasts. In *BuildSys '10*.
- [16] A. Rice, S. Hay, and D. Ryder-Cook. A limited-data model of building energy consumption. In *BuildSys '10*.
- [17] T. Teixeira and A. Savvides. Lightweight people counting and localizing in indoor spaces using camera sensor nodes. *ICDSC*, 2007.