Implicit Sensing of Building Occupancy Count with Information and Communication Technology Data Sets

Bianca Howard[a,*], Salvador Acha[b], Nilay Shah[b], John Polak[c]

[a]*School of Architecture, Building, and Civil Engineering, Loughborough University, United Kingdom*
[b]*Department of Chemical Engineering, Imperial College London, United Kingdom*
[c]*Department of Civil and Environmental Engineering, Imperial College London, United Kingdom*

Abstract

Occupancy count, i.e., the number of people in a space or building, is becoming an increasingly important measurement to model, predict, and minimize operational energy consumption. Explicit, hardware-based, occupancy counters have been proposed but wide scale adoption is limited due to the cost and invasiveness of system implementation. As an alternative approach, researchers propose using data from existing information and communication technology (ICT) systems to infer occupancy counts.

In the reported work, three different data streams, security access data, wireless connectivity data, and computer activity data, from ICT systems in a medium sized office building were collected and compared to the counts of a commercially available occupancy counter over 59 working days. The occupancy counts from the ICT systems are compared to the commercial counter with and without calibration to determine the ability of the data sets to measure occupancy. Various transformations were explored as calibration techniques for the ICT data sets. Training sets of 24, 48, and 120 hours were employed to determine how long an external calibration system would need to be installed.

The analysis found that calibration is required to provide accurate counts. While each ICT data set provides similar magnitudes and time series behavior, incorporating all three data streams in a two layer neural network with 1 week of training data provides the most accurate estimates against 5 performance metrics. Whilst 1 week of data provides the best results, 24 hours is sufficient to develop similar levels of performance.

1. Introduction

The following subsections introduce the motivation and need for collecting occupancy count data, the various methods used to measure occupancy, and the aims and significance of this work.

1.1. Motivation

Countries across the world have made pledges to reduce greenhouse gas emissions to reduce the impact of climate change. The IPCCs fifth assessment report identifies energy demand reduction, through increased energy efficiency and behavioral change, as key to achieving mitigation targets [1]. For the buildings sector, this translates into policies to encourage the uptake of energy efficient devices, systems, and operational strategies for the provision of space heating, water heating, space cooling, lighting, appliances, and cooking demands. With the advancement of digital technologies, energy efficient operation and behavior change can be implemented through intelligent automation of systems, especially for the provision of heating cooling and ventilation (HVAC) systems in commercial buildings.

---

*Corresponding author.
Email address: `b.howard@lboro.ac.uk` (Bianca Howard)

Intelligent building systems attempt to dynamically adjust the buildings operation to minimize energy consumption considering variable conditions without sacrificing occupant comfort. Traditional control systems have minimal ability to alter their operation in response to changing occupancy patterns and weather conditions. At most, these systems change operational conditions based on fixed times of the day, e.g. time clocks, or through simple relationships between outdoor and desired indoor temperatures, i.e. set backs [2]. The new generation of automated building control systems will be able to sense the current occupancy and other drivers of energy demand, forecast their likely future values, and operate the system to maintain thermal comfort with minimal energy consumption. Key to the operation of intelligent building systems is the ability to sense when a building or space is being used.

In current buildings, some measures of occupancy sensing have been incorporated, specifically for lighting and ventilation systems. Many lighting systems sense occupant presence through passive infrared (PIR) sensors. These sensors have been extensively researched and implemented in buildings [3]. PIR sensors detect a change in the amount of thermal radiation that falls on its surface through occupant movement. The lighting system will turn on the lights in the space if movement is detected and shut off the systems if no movement is detected for a specific time period. These systems can achieve up to 86% reduction in building lighting energy, where the magnitude of the savings depend on the current lighting operations and building use patterns [4].

Sensors detecting occupant activity use have also been deployed to operate commercial ventilation systems. As people exhale carbon dioxide into a space, the carbon dioxide (CO_2) concentration can act as a measure of the number of people in the space but also a proxy for air quality. Therefore many demand-controlled ventilation systems set targets for CO_2 concentrations. The current CO_2 concentration is measured in the building and the ventilation system responds to ensure the CO_2 concentration remains below this value. Demand-controlled ventilation has achieved savings of up to 50% of fan energy consumption and 30% of total HVAC energy consumption as compared against strategies that do not consider occupancy levels [5].

However more energy savings can be achieved if occupancy driven control is incorporated in the operation of heating and cooling systems as well. Model predictive control (MPC) strategies have the potential to enable more energy savings based on developing optimal strategies of building operation [6]. MPC strategies for heating and cooling systems use information about the building thermal dynamics, current weather, and internal conditions to determine the optimal set of control actions to achieve an objective function, typically to reduce energy or energy costs. These systems have shown to use information about upcoming occupancy presence to drastically reduce energy consumption during unoccupied periods, similar to optimal start/stop routines, and take advantage of free heating from internal gains for occupants [7]. As reviewed in [6], several researchers conducting experimental research have incorporated information about occupancy within their algorithms to achieve better operating HVAC systems.

Bengea et al. [8] evaluated the ability of a model predictive control strategy to reduce energy consumption in a commercial building. In their experimental set up they measured occupancy using a combination of motion, CO_2, and bidirectional people counters [9]. These data were used to calculate the internal gains and occupied periods for multiple zones in a commercial building. Through the MPC framework, they obtained 30% energy savings. Aswami et al. [10] have developed an online learning mpc to reduce electricity consumption of an air conditioner serving a computer lab. Instead of measuring occupancy directly, the internal loads were identified and forecast from a model of building thermal dynamics and measured temperatures. Occupancy periods were identified by the time of day. Through this MPC framework, savings of 30-70% in electricity consumed for cooling were realized. Dong et al. [11] developed a non-linear MPC to control a home heating system, with solar thermal heating and radiant floors, and a cooling system with fan coil units. The nonlinear MPC considered estimates of building occupancy count and occupancy duration using a combination of CO_2, acoustic, motion and lighting change sensors. They found a 30% reduction in energy consumption during the heating season and 17% reduction in the cooling season. In each of the previous examples, an understanding of building occupancy, or internal loads from building occupancy, was needed. However the researcher, who can be seen as early adopters, used many methods to estimate building occupancy as a single method for sensing building occupancy count has not yet been developed.

1.2. Techniques for Measuring Occupancy

The lack of standard procedures and methods for measuring building occupancy has received a separate treatment in the literature. There have been several researchers that have reviewed the current state of measuring occupancy in buildings. Recent reviews include [12, 13, 14]. As many reviews in the literature exist, the following sections provide a brief overview of techniques and is not meant to be comprehensive. The following paragraphs describe the use of explicit sensing techniques, implicit sensing techniques, the scale of current occupancy sensing studies in the literature, and the calls made for future research in the area.

Explicit sensing systems are those that introduce new hardware to sense the presence of building occupants. In the literature, a variety of mechanisms have been used to detect occupancy counts. Pressure sensitive floor systems have been proposed for capturing the presence of individuals by measuring the pressure throughout an instrumented floor space. When placed together in a network these systems can also provide occupancy count and localization. Various versions have been proposed from dense pressure sensitive tiles to smart carpets [15]. The main technical challenge with this technology is to determine the density of the sensing network and appropriate algorithm to accurately differentiate the pressure profiles of various individuals. The main criticisms of this approach are the intrusiveness of implementation, the high-density network required for accuracy and the difficulty in counting with large numbers of people [16]. Electric field sensors, where capacitors are placed throughout the space to detect the changes in electric field caused by people, ultrasonic and vibration sensors while using different sensing modalities suffer from similar drawbacks as the pressure sensitive floor systems.

Other, slightly less infrastructure heavy techniques, include camera-based systems and wearable devices. In camera-based systems, images are taken and occupants are identified from the frames through object tracking software [17, 18, 19]. The hardware for this approach is widely available although there are technical challenges in developing algorithms that can accurately detect and count people in varied settings. Even in the face of these challenges, these systems have been shown to provide detection rates of up to 97%. However with camera based systems, one must take care to address privacy concerns as occupants are being visually monitored.

With systems employing wearable devices, occupants are asked to attach a beacon or receiver to their person that actively or passively interacts with the environment. Typically this is in the form of mobile devices, such as smart phones and watches, or specially developed RFID tags [20, 21]. As identification and location of a device is core to the development of wireless information and communication devices, these systems can report high accuracies. Overall the major criticism of this approach for general sensing applications is that each occupant must wear the device at all times and if the device requires power it must be kept charged [22].

In contrast to explicit sensing techniques, implicit techniques attempt to infer occupancy using data currently collected by the building for other purposes. These sensing methods assume that building occupants alter the environment that they occupy, and these disturbances can be captured through the various sensing mediums. Implicit occupancy techniques have been classified as Tier I, II and III by Melfi et al [23]. Tier I requires no additional infrastructure, just collection and processing. Tier II requires additional software to access the data available in existing infrastructure. Tier III requires the additional software and hardware to estimate occupancy. Tier I and II sensing systems are highly desirable as they do not require any additional infrastructure, reducing the overall cost of the sensing apparatus and potentially increasing the scalability of the sensing methods. Shen et al. in a recent review of 50 implicit sensing systems reported systems using data streams collected from internet traffic, outbound phone calls, access badges, Wi-Fi systems, online messaging services, keyboard movements, and webcam usage to detect occupancy [12].

The studies explicitly considering occupancy count reported high accuracies with a variety of sensing data streams and data fusion techniques. The same CO_2 sensors used for demand controlled ventilation can also be used to estimate occupancy counts. These systems estimate the number of people in a space with carbon dioxide sensors with inverse modeling through a dynamic mass balance equation. The inverse modelling technique requires assumptions about the amount of carbon dioxide that is released by the individuals, carbon dioxide levels in the local outdoor climate and infiltration rates. Studies have shown that this system

can provide counts with accuracy between 81% [24] and 94% [25]. For large scale occupancy detection, an array of carbon dioxide sensors would need to be implemented and the potentially time varying constants for each sensor estimated. Ekwevugbe et al. evaluated the ability of environmental data, using added sensors measuring sound, VOC, lighting, temperature, relative humidity, and CO_2 measurements, to reflect occupancy count for an office of 6 people [26]. The data were combined and transformed using a neural network to reflect occupancy counts. Considering various feature combinations, they reported accuracies up to 74% and root mean squared error (RMSE) values as low as 1 person. Whereas Christensen et al. [22] reviewed the ability internet traffic, outbound phone calls, access badges and computer activity to indicate occupancy in academic labs. They reported DHCP leases could reflect building occupancy with accuracy of 89%.

Implicit techniques have the ability to reflect occupancy however developing a consistent methodology is difficult as different buildings have different availability of sensors and the techniques for combining data streams is quite bespoke. Further while many studies have been performed to assess the ability of a variety of sensing mechanisms in buildings, the majority of studies have only reflected small scale experiments which is reflected in the their methodological approach.

A review by Shen et al.[12] identified 27 studies that have validated the ability of their occupancy count sensing technology against a ground truth. Twelve studies used direct manual counting of the occupants in the scene and remainder used camera systems. Even within the studies that used cameras to establish a ground truth, the majority manually counted the number of occupants in the scene. As the majority of studies consider small numbers of occupants (<15), manual counting was feasible. As implicit occupancy is in its early stages, developing ground truth techniques that apply at scale are needed to allow for the evaluation of sensing mechanisms in typical buildings.

The vast amount of sensing approaches and experiments performed in this area led to the formation of an International Energy Agency (IEA) Energy in Buildings and Communities (EBC) Annex to develop a standardized definition and approach to understanding occupants and their influence on building energy consumption. The final report of the IEA EBC Annex 66 [14] entitled the "Definition and Simulation of Occupant Behavoir in Buildings" has reviewed the current techniques and challenges to understanding occupancy behavoir in buildings. One of the conclusions the expert group made upon reviewing 14 state of the art occupancy sensing techniques was that there is no one technology that can detect both presence and occupancy count in a cost-effective way with high accuracy. Further as most studies are on a small scale, one of the highlighted areas of future research is the need to determine reliable and affordable ways to collect large-scale occupant behavoir data.

1.3. Aims and Significance

In summary, incorporating occupancy data in building heating ventilation and air conditioning systems can lead to significant energy reductions as compared to standard approaches. However there is a need to determine a scalable and cost effective technique to measure occupancy count and presence to deliver those savings. Previous literature indicates two approaches for sensing occupancy, explicit and implicit sensing techniques. Of the two, implicit sensing is a potential way forward as previous studies indicate that implicit occupancy sensing can provide accurate and cost-effective occupancy counts with potentially no need for additional monitoring infrastructure as the building increases in size. However, before definitive statements can be made, there is a need to test the effectiveness of implicit techniques at scale.

The aim of this work is to address this gap by assessing the ability of implicit occupancy sensing techniques to detect occupancy count for a large office building. Data reflecting desktop computer activity, security access and mobile Wi-Fi connectivity for a 350-person office building was collected over 59 weekdays and compared the occupancy counts to those reported by a commercially available stereo-vision based occupancy counter. Further, various relationships were evaluated to calibrate each individual data set as well as combinations of the ICT data sets. Various training and testing data sets lengths were evaluated to determine the trade-off between calibration period and accuracy. The ICT data sets considered for this work are considered Tier I, as all software needed to collect the data was already installed in the current ICT network.

This paper reports the first building scale comparison of explicit and Tier I implicit occupancy sensors, contributing to the building occupancy measurement field by structuring a measurement methodology, ex-

ploring various calibration procedures, and demonstrating the ability of implicit sensing to reflect occupancy count at scale.

1.4. Organisation

The remainder of this paper is organized as follows: section 2 describes the case study building and location of sensing devices, the data collection methods, model structures used to calibrate the data sets as well as the performance metrics used; section 3 reports the results comparing the ICT datasets, calibrated and uncalibrated, to the counts from a commercial occupancy counter; Section 4 provides a discussion of the results; Section 5 discusses the limitations of the work and areas for future research and Section 6 ends the paper with a statement of the main conclusions.

2. Methodology

The following sections describe the case study building, the technology behind the commercial occupancy counter, how occupancy estimates were derived from each of the ICT data sets, and the metrics used to evaluate the implicit sensing performance.

2.1. Case Study Building

The case study building is the administrative hub of a University. There are 350 people assigned working space in both open plan and closed offices. The building also contains three meeting rooms capable of accommodating 60 people in aggregate. As it is the administrative hub, many employees enter and exit the building throughout the day to go to other meetings throughout the campus. On the mezzanine level, the building also houses secure space to store bicycles for those commuting to work.

Access is gained to the building through two entry/exit points: the main entrance and the rear bike rack entrance. The main entrance is equipped with entry gates that release upon the presence an ID card with appropriate permissions. For the rear bike rack entrance, presenting an ID card to the reader unlocks a set of doors. The building's wireless network is equipped with eight wireless access points, two on each floor of the building. Within the building, there are 321 desktop computers equipped with a power management software to track their power consuming activities. Specifically the software reports how many minutes a computer was active, the computer was on, and the monitor was on. The computers are located through out the building in the open plan office space as well as closed offices.

The commercial occupancy counters were placed above each entry/exit point to the building. They were installed to ensure each person entering the building at various angles would be captured in the field of vision. Figure 1 shows the building layout and the positions of the occupancy counters, security access points, and wireless access points.

2.2. Data Collection and Repair

The following sections describe the data collection approaches for the commercial counter and the ICT data sets.

2.2.1. Commercial Occupancy Counter

Many organizations have an interest in understanding occupancy within their spaces, especially in the retail sector. Therefore several companies have developed sensors to count and track visitors based on various technologies such as thermal imaging, infrared beams, and 3D stereo video. The commercial counters used in this work are 3D stereo video cameras equipped with proprietary image processing software to detect the height, direction, mass and velocity of people. As the cameras have a limited field of vision, to measure building occupancy they must be placed at each entry and exit point to ensure all traffic in and out of the building is recorded. Occupancy is then determined by a cumulative sum of the net flows of traffic.

Overall building occupancy is determined by

Figure 1: Case Study Building Layout and Location of Monitoring Points
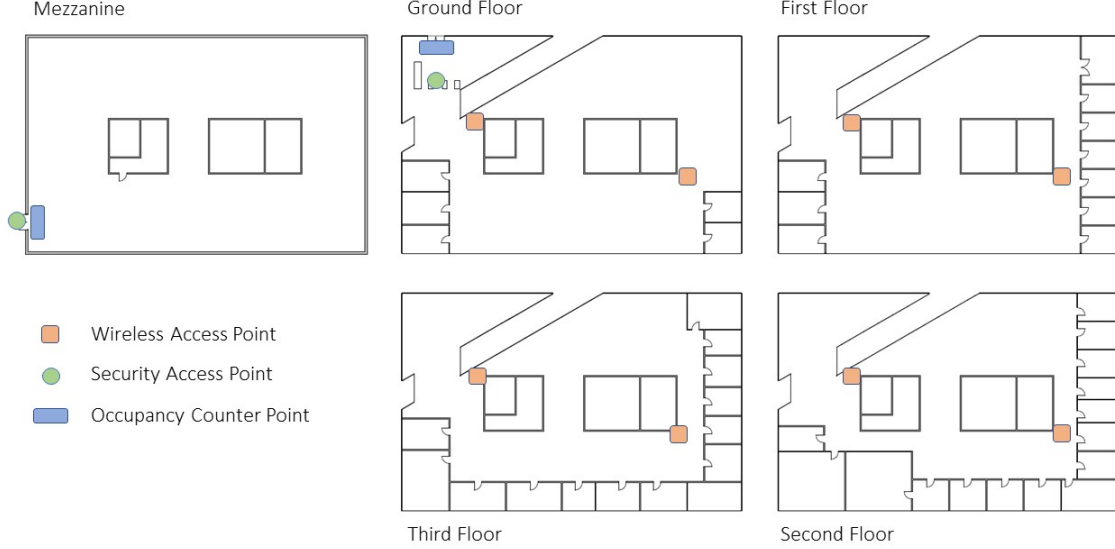
$$O_t^c = \sum_{s=1}^{2} \sum_{k=1}^{t} f_{k,s}^c - \sum_{s=1}^{2} \sum_{k=1}^{t} r_{k,s}^c \qquad (1)$$

where O_t^c is the occupancy count measured by the commercial counter at the end of time period t, $f_{k,s}^c$ is the number of people measured entering the building within time period k from entry s, and $r_{k,s}^c$ is the number of people measured exiting within time period k from entry s. The software accompanying the commercial counter allows for reporting of occupancy at the end of 15 minute intervals $[k]$. To allow for comparison with the ICT data sets, the occupancy was reported at hourly intervals $[t]$.

Occupancy counters based on counting the number of occupants crossing thresholds have shown to be prone to accumulation of errors [27, 28, 29]. If a person is not detected leaving a building once having entered, it results in an ongoing increase in the count of occupants. Two approaches have been proposed to adjust for this error [28]: assuming the building is empty at end of the day and incorporating an algorithm to adjust the counts. Both approaches were applied for the occupancy counter deployed in this work.

The assumption that the building is empty at the end of the day is reasonable as the case study building is closed at 11pm. Security guards check the building after closing time to ensure the building is empty. Second the algorithm used to remove the bias does this by assuming fewer people left the building during the day, and that this occurred during high flow periods. However, the full details of the approach are commercial intellectual property. To achieve the occupancy counts, counters were placed at the main entrance and the rear bike rack entrance by the commercial installers. Over the course of a week the counter was physically adjusted and the algorithm modified based on the commercial installers viewing the raw video footage. After this adjustment period, the installers verified the reported accuracy of 95%.

2.2.2. Computer Activity Data

For each desktop computer in the case study building, for power management purposes, the computer activity data is tracked. Specifically the system logs three different states: the minutes in each hour the computer was active "computer active", the computer was on "computer on", and the monitor was on "monitor on". Occupancy from this data set was defined with the "computer active" state as,

$$O_t^a = \frac{\sum_{n=1}^{N} m_{n,t}}{60} \tag{2}$$

where O_t^a is the occupancy count derived from the computer activity data at time t, $m_{n,t}$ is the number of activity minutes for desktop n in time t and N is the total number of desktop computers (321). This expression provides a fractional value of occupancy as it is assumed that each minute in an hour represents a fraction of a single individuals presence. This approach for counting occupants would only capture the employees actively engaged with computer based work. Those not working on a computer or visitors to the building would fundamentally not be captured through this approach.

2.2.3. Wi-Fi Connected Mobile Devices

The case study building has 8 wireless access points, 2 on each level, that allow occupants to connect to the wireless network. For a device to send data through the network, it must first undergo two steps: an authentication and an association. The authentication ensures that the mobile device and wireless network are compatible while the association allocates a wireless access point to a device.

Counts of the number of devices authenticated and associated with the network were sampled hourly. Preliminary data analysis found that there is minimal difference between the number of devices connected at the two steps, although there are consistently more devices associated with the network than authenticated. In the current analysis, the hourly counts of associated devices were used as a measure of occupancy, assuming that the devices associated with the network are still physically present within the building.

This approach assumes that each occupant in the building is carrying a mobile device currently communicating with the wireless system. This would not capture people not carrying a device or those where their device is powered down. Further this approach would over estimate occupancy if devices were left in the building.

2.2.4. Security Access Data

Access to the building is provided by presenting an ID card with appropriate permissions. The access control system records each entry and exit in a data log which contains the direction of entry (IN or OUT), a unique ID for each person, the campus building to which they are assigned, and the location of their office with in that building. As with the commercial counter, the access data provided the flows of people in and out of the building. Therefore the building occupancy can be derived from the following equation

$$O_t^s = \sum_{p=1}^{2} \sum_{k=1}^{t} f_{k,p}^s - \sum_{p=1}^{2} \sum_{k=1}^{t} r_{k,p}^s \tag{3}$$

where O_t^s is the occupancy count as measured from the security access data at time t, $f_{k,p}^s$ is the number of people entering during time period k from access point p, and $r_{k,p}^s$ is the number of people exiting during time k from access point p.

Upon initial data exploration, it was clear that there were missing data points, in that there were times when people would be registered entering the building but not exiting or vice versa. After observation of each entrance and exit, these missing data points can be attributed to the tailgating phenomena, where a single person may present an ID but several people may follow behind, resulting in their entry or exit not being logged. To account for the missing data, a repair methodology was used where missing values were imputed using historical empirical distributions of arrival and departure times of each individual. In the analysis this repair was performed at the end of the day. As implemented, this would prevent the real time use of the access control data. However these estimates could be used to estimate occupancy in real-time updating for incoming data. This would most likely lead to less accurate occupancy estimates in addition to the current imputation strategy which is a limitation of the work. Full details of the imputation strategy can be found in previous works [30].

2.3. Calibration and Adjustment Methodology

In many sensors calibration and adjustment may be needed to enable the ICT data sets to reflect accuracy estimates of occupancy count. In this case the aim is to determine a function that will map the implicitly measured occupancy counts to the explicitly measured counts. To determine the nature of the relationship, many model structures were considered and compared. For each model structure, parameters are identified on a training data set and evaluated on a testing data set. In this case, the length of the training set is of specific interest. It represents the amount of time required for calibration, or alternatively, how long an external system would need to be installed to estimate the model parameters. Ultimately, it is favorable to have a short calibration period, or training data set. Therefore, relatively small training sets are evaluated, when compared against traditional data mining approaches. With 59 days of hourly data collected, the ability of training sets of 24, 48, and 120 hours were evaluated, representing between 1.7% and 8.5% of the entire data set. Several sets of the data are considered for training and testing. The training data sets are defined as a k-fold cross validation, where the value of k is chosen such that each training set contains the desired number of samples for validation.

2.3.1. Linear Regression

The simplest adjustment that can be made is to multiply the occupancy count by a single factor and add a bias, in other words a linear relationship. Ordinary least squares regression can be used to estimate the values of the factors. Specifically

$$y_t = \beta_0 + \beta_1 x_t + \epsilon \tag{4}$$

where y_t is the occupancy counter as measured by the commercial counter in time step t, x_t is the occupancy count as measured by the ICT data set in time step t, β_1 is a constant coefficient that change the amplitude of the measurement, β_0 is a constant coefficient representing an offset, and ϵ represents the noise in the data assumed to be normally distributed with mean 0 and variance σ^2. The coefficients of the model are estimated using the maximum likelihood estimator.

Each ICT data set, computer activity, security access, and WiFi connectivity data, are individually calibrated with univariate linear regressions. The models are identified in the remainder of the paper as the "linact", "linacc", and the "linwif" representing the linear regressions with the computer activity data, security access data and Wifi connectivity data, respectively.

However it is also possible for each data stream to be combined to provide an aggregate estimate. This is done in two ways: a multiple linear regression and data fusion. In the multiple linear regression, denoted "linmul", each data set is considered in a single model specifically,

$$y_t = \beta_0 + \sum_{i=1}^{3} \beta_i x_i + \epsilon \tag{5}$$

where x_i represents one of the three ICT data streams.

2.3.2. Data Fusion

For the data fusion approach denoted "fusion", the forecasts from each individual model were combined to develop a single measurement. It has been shown [31, 32] in the presence of uncorrelated errors that the optimal weights for combining separate forecasts are a function of the variance of each model.

Specifically the forecast combination estimate is defined as

$$\mu_f = \sigma_f^2 \sum_{i=1}^{3} \frac{\mu_i}{\sigma_i^2} \tag{6}$$

where

$$\frac{1}{\sigma_f^2} = \sum_{i=1}^{3} \frac{1}{\sigma_i^2} \tag{7}$$

where μ_f is the data fusion based occupancy count , μ_i is the occupancy count as predicted with ICT data set i, and σ_i^2 is the variance of the residuals as predicted with ICT data set i. An assumption for this approach is that the errors are Gaussian which by observation is not strictly the case for the data. Therefore these weights may not be optimal for the current forecasts and is a limitation of this implementation of the approach. The approach was still considered for its simplicity.

2.3.3. Artificial Neural Network

To develop models capable of considering nonlinear and time lag relationships, neural networks were used. These models were chosen as sufficiently complex neural networks have the ability to replicate any function. The approach taken in this work was to use neural network models as indicators of nonlinear or time lag behavoir. If increased performance was shown, then future work would be needed determine the exact form of this relationship. For the analysis, three different types of neural networks were considered: a standard neural network, a recurrent neural network, and a long short term memory neural network.

A neural network consists of a composition of neurons that transform inputs to outputs. The general form of a single neuron is

$$y = f\Big(\sum_{i \in I} w_i x_i + b\Big) \tag{8}$$

where y is the neuron output, x_i is the neuron input, w_i are the weights applied to each input, b is the constant bias, $f()$ is a function representing the activation function and I is the set of all inputs. Neural networks can have multiple neurons in many configurations. The most common configuration is a series of layers, where the output of one layer of the neural network composed of an arbitrary number of neurons is used as the input to a second layer of neurons. The current work considered a neural network with a single layer network with single neuron (NN1) and a two layer network with three neurons in the first layer and one neuron in the second layer (NN3).

The NN1 model has the same formulation as equation 8 where the inputs, x_i, are the three ICT data streams and the output is the estimated occupancy count. The NN3 model has the following formulation

$$y = f^o\Big(\sum_{d \in H} w_d h_d + b_o\Big) \tag{9}$$

$$h_d = f^d\Big(\sum_{i \in I} w_i x_i + b_d\Big) \tag{10}$$

where f^o is the activation function for the second layer (single node), f^d is the activation function for the first layer d (3 nodes), w^d is the weight of output d from layer 1, w_i is the weight for input i, x_i is input i, b_o is the constant bias for the output node, b_d is the constant bias for node d in the first layer and O is the set of the outputs from the first layer (3).

2.3.4. Recurrent Neural Network

A recurrent neural network allows for consideration of a time series, acknowledging that information in the recent past can be used to predict the future value. The model structure allows information from an estimate for a previous time step to be available for estimating the next time step. A node of a recurrent neural network has the form

$$y_{o,t+1} = f\Big(\sum_{i \in I} w_i x_{i,t} + \sum_{o \in O} r_o y_{o,t} + b\Big) \tag{11}$$

where $y_{o,t+1}$ is the output of node o at time t, $y_{o,t}$ is the output of node o at time t, r_o is the weight assigned to the recurrent input y_t. In this form, the prediction from the previous time step is used to inform the prediction from the next time step. The current work considered a recurrent neural network with a single layer network with a single neuron (RNN1) and a two layer network with three neurons in the first layer and one neuron in the second layer (RNN3).

2.3.5. Long Short Term Memory Network

As time steps forward, less information from the beginning of the time series is used to inform that latter part of the time series. This means that a recurrent neural network can only remember information from the immediate past. To incorporate longer times series behavior, long short term memory (LSTM) networks have been developed and are considered as a calibration strategy.

LSTM neural networks incorporate a model structure that allows the model to decide how much information to keep and how much to forget over time. A neuron in a LSTM neural network has the form

$$y_{o,t+1} = f^4\left(i_t^4\right) * tanh\left(c_{t+1}\right) \tag{12}$$

$$c_{t+1} = c_t f^1\left(i_t^1\right) + f^2\left(i_t^2\right) * f^3\left(i_t^3\right) \tag{13}$$

$$i_t^g = \sum_{i \in I} \beta_i^g x_{i,t} + \sum_{o \in O} a_o^g y_{o,t} + b_g \tag{14}$$

where $y_{o,t+1}$ is the output of node o at time t+1, $x_{i,t}$ is the input i at time t, $y_{o,t}$ is the output of node o at time t, β_i^g is the weight for input i of function g, a_i^g is the recurrent weight for output o of function g, b_g is the bias factor of function g, and $f^g()$ is the activation function g. The current work considered a LSTM network with a single layer network with a single neuron (LSTM1) and a two layer network with three neurons in the first layer and one neuron in the second layer (LSTM3).

Each of the neural networks was solved using backpropagation with stochastic gradient descent implemented with the Keras python API. The loss function was the root mean squared error. The weights were estimated considering 10,000 epochs with an additional minimal change stopping criterion equal to 1e-6.

For all of the neural networks considered, linear activation functions were used. In preliminary analysis, the linear activation functions yielded the best performance against hyberbolic tangents, sigmoid, and rectified linear units. This means that the NN1 model has a linear model structure whilst the remaining neural networks have a nonlinear structure.

2.4. Performance Metrics

The performance metrics are used to determine how well the implicitly sensed occupancy count reflects the explicitly sensed occupancy count from the commercial counter in the calibrated and uncalibrated settings. The metrics used attempt to provide a sense of the overall fit, ability to measure arrival and departure times,as well as peak occupancy and the time at which it occurs. All comparisons were made for weekdays only. Weekends were not included as the occupancy was close to zero throughout the period.

Two metrics were used: the mean absolute error (MAE) and the accuracy. MAE is used to determine the difference between measures of occupancy count and accuracy is used to denote the ability of implicit sensors to determine the timing of events. MAE was chosen as a metric as it allows for easy interpretation of results. In this scenario there is no need to penalise larger errors as is done with the root mean square error (RMSE), a commonly used metric.

The mean absolute error (MAE) is used as a metric to evaluate the implicit occupancy sensor to reflect occupancy count overall and to reflect the peak occupancy in a 24 hour period. The MAE is defined as

$$MAE = \frac{1}{n} \sum_{n=1}^{N} \mid y_n - \hat{y}_n \mid \tag{15}$$

were y_n is the occupancy as measured by the implicit occupancy sensor at point n, $\hat{y}_n$ is the occupancy count measured by the explicit occupancy sensor at point n, and N is the total number of samples in the set. MAE_o is used to denote when N includes all 59 days of hourly data in the uncalibrated set (and the full testing data set in the calibrated case), and n is equivalent to the hour of measurement. MAE_p is used to denote when y_n and $\hat{y}_n$ represents the peak occupancy occurring over a single day (24 hours) for the implicit and explicit occupancy sensor, respectively. N represents the number of days in the data set

Figure 2: Occupancy Count as measured from Commercial counter and ICT datasets for the week of January 30th 2017. access: Security access data; wifi: WiFi connectivity data; activity: computer activity data; counter: commercial occupancy counter
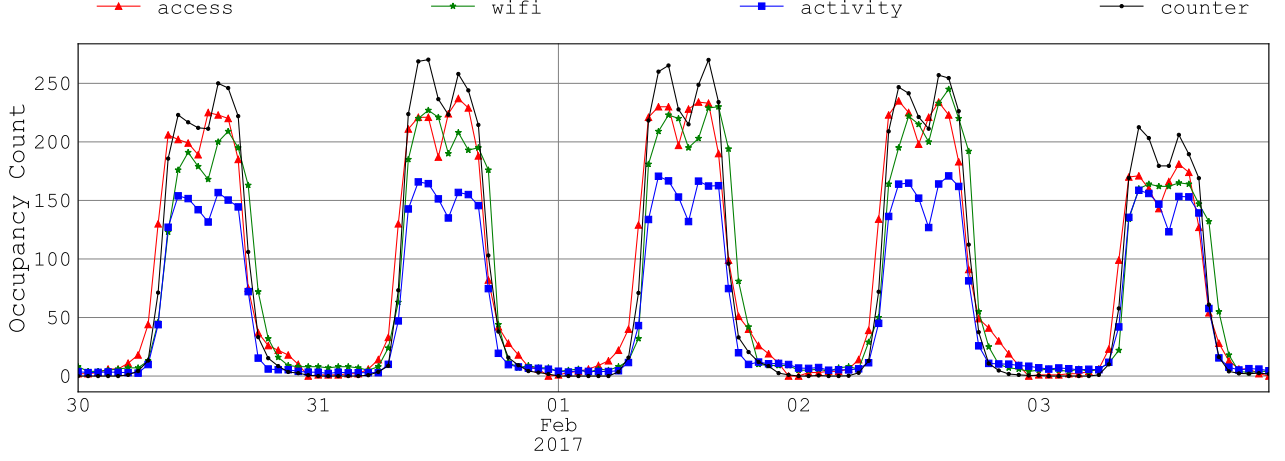
and n denotes the specific day. For the MAE_p metric, the timing of the events in the implicit and explicit occupancy counts may not be the same, as this metric only considers the magnitude of the peak occupancy, not when it occurs.

The accuracy metrics are determined by dividing the total number of correct timings by the total number of events. Three different accuracy metrics are reported: the arrival time accuracy (ATA), the departure time accuracy (DTA) and the peak occupancy time accuracy (POTA). Given the large building occupancy, the time of arrival is defined as the time when the occupancy first surpasses 5% of the maximum occupancy and the time of departure is defined as the time when the occupancy, after first surpassing 5% occupancy, subsequently drops below 5% occupancy. The peak occupancy time compares the daily peak occupancy time from the implicit sensor to that by the explicit sensor.

The commercial occupancy counter has a reported accuracy of 95%. To provide a consistent metric for evaluation, the MAE_o, MAE_p, POTA, ATA, ans DTA were calculated for the upper and lower bounds of the accuracy for reported occupancy counter. Therefore the metric values that fall within the error of the commercial counter are $MAE_o = 4.2$, $MAE_p = 12.4$, POTA = 1, ATA = 1, and DTA = 0.90-0.92. The departure time accuracy is reported as a range as the upper and lower bounds lead to different accuracy figures.

3. Results & Discussion

3.1. Uncalibrated Occupancy Count

Table 1 depicts the performance metrics for the uncalibrated data sets. In comparison to the commercial occupancy counter, each ICT data set has a large error in both the overall MAE and the peak MAE, with the WiFi connectivity and security access data streams providing closer estimates of occupancy count. This can also be seen in time series data depicted in Figure 2 where the activity data results in a consistently lower occupancy count. This difference is highlighted in the peak mean absolute error metric, where the activity data underestimates the peak by 85 people on average in comparison for 31 and 35 people for the security access and WiFi connectivity data streams, respectively.

However with respect to the timing of events, the activity data out performs the other data streams. The activity data estimates the arrival time with 97% accuracy and departure time with an accuracy of 68%. The security access data whilst providing the lowest MAE does not perform well in describing the timing of activities with an arrival and departure time accuracy of 3% and 12% respectively. This could perhaps be due to the imputation strategy used to repair the data set where timings of arrivals and departures are

11

Table 1: Performance Metrics for Each Uncalibrated ICT Data Set and each calibrated model. For the calibrated models, the metrics are the median values considering 59 calibration periods. MAE_o: Overall Median Absolute Error; MAE_p: Peak Median Absolute Error; POTA: Peak occupancy time accuracy; ATA: Arrival Time Accuracy; DTA: Departure TIme Accuracy

ICT Data Set or Model	MAE_o	MAE_p	POTA	ATA	DTA
Commercial Counter Error	4.2	12.4	100%	100%	90%
Uncalibrated					
Access	17.3	30.6	31%	3%	12%
WiFi	19.8	35.4	34%	59%	36%
Activity	29.3	85.4	39%	97%	68%
Calibrated					
fusion	7.2	15.2	41%	76%	69%
linmul	7.3	14.4	37%	69%	66%
linact	8.4	17.9	40%	97%	69%
NN3	10.1	17.3	42%	95%	67%
NN1	12.6	14.6	39%	12%	28%
RNN	14.7	13.8	26%	12%	47%
linacc	15.4	13.7	31%	10%	25%
LSTM	15.7	14.3	36%	79%	5%
linwif	16.0	20.0	34%	76%	42%
LSTM3	17.5	15.9	36%	39%	3%
RNN3	21.4	21.2	26%	53%	34%

estimated. All data sets have around a 33% accuracy of determining when the peak occurs. The occupancy count for this office building displays a dip in occupancy over the lunch time hours, leading to two potential peaks, a pre-lunch peak or a post-lunch peak. Estimating the peak occupancy time accurately requires the data stream to accurately measure the count at both times or else one would identify the peak time at the wrong time. An example of this behavior can be seen in Figure 2. On the 31st of January, the security access control data estimates the highest daily occupancy count occurring post lunch where as the commercial counter registers the maximum daily occupancy at the pre-lunch peak.

Viewing the time series data in Figure 2, one can also view some qualitative differences in the shape of each data stream. The activity data has a consistent profile with a distinct dip at lunch time. The WiFi connectivity data stream sometimes reflects the shape of the occupancy as reported by the commercial counter but appears to lose the signal, apparent in Figure 2 on the 31st of January and 3rd of February. On January 31st, the wireless data fails to capture the afternoon peak and on February 3rd fails to capture the afternoon dip. The WiFi connectivity data also appears to possibly have a slight delay as indicated by the consistently higher occupancy counts during the day's evening hours. The security access data also shows a distinct lunch time dip with magnitudes similar to the commercial counter. It does however exhibit higher occupancy counts during the late evening hours and early morning hours leading to the low arrival and departure time accuracy.

3.2. Calibrated Occupancy Counts

Several model structures were evaluated for their ability to calibrate each of the data streams. Training sets of 24 hours, 48 hours, and 120 hours were evaluated to determine how the accuracy would change with longer training sets, enabling one to understand how long an external calibration system would need to be placed. The "Performance of Model Structures" section below reports the performance metrics considering training data sets of 24 hours. The "Effect of Training Set Length" section describes the changes in the calibration metrics for the longer training data sets.

3.2.1. Performance of Model Structures

Considering the overall median MAE, all model structures once calibrated, with the exception of RNN3 and LSTM3, led to a lower MAE than with the best uncalibrated data set as shown in Table 1. The median peak MAE for all calibrated models was lower than the best uncalibrated models. However none of the model structures led to MAE_o or MAE_p that were less than those of the commercial counter or POTA, ATA, DTA that were as high as those of the commercial counter. This indicates that even with calibration the ICT data sets were not able to estimate occupancy as accurate as the commercial occupancy counter.

Of all of the model structure considered, the "fusion" model led to the lowest MAE_o and the univariate regression with the security access data, "linacc" led to the lowest MAE_p. The univariate regression with the computer activity data, "linact", resulted in the highest arrival and departure time accuracy. The 2 layer neural network (NN3) resulted the highest peak occupancy time accuracy. There was not one model structure that performed best across all metrics considered. More detailed results for each of the main model structure types are discussed below.

3.2.1.1 Univariate Linear Regressions

The univariate linear regressions for each data set had varied reductions in the MAE. The security access data, which had the lowest uncalibrated MAE of 17.3 people, through calibration the MAE was reduced to 15.4 people. The WiFi connectivity data reduced from an overall MAE of 19.8 to a median MAE of 16 people. However the activity data which had the worst uncalibrated MAE, through the linear regression, resulted in a MAE of 8.4 people. The MAE for the peak occupancy is reduced for all data sets, with the activity data having the most significant change from a mean absolute error of 85 people to a median MAE 17.9 people. Overall, the linear regression with the security access data ["linact"] was, out of all of the model structures evaluated, best able to estimate the peak occupancy count with a median MAE of 13.7.

Figure 3 depicts the results of each of the linear regressions for the week staring Monday January 30th, 2017. As compared to Figure 2, each of the data streams better aligns with the values reported by

13

the commercial occupancy counter. The calibration for the activity data however leads to a large over estimation of the occupancy on Friday. In viewing the results for the remainder of the year, this phenomena occurs frequently indicating there may be a different relationship between occupancy and computer activity on Fridays. As the model structure is simply a linear transformation, the long tails of the security access data and the delayed count of the WiFi connectivity data previous discussed are still present. The metrics for arrival time, departure time, and peak occupancy time only slightly improve for the calibrated data sets.

3.2.1.2 Multiple Linear Regression and Data Fusion

The fusion and the multiple linear regression had the lowest median MAE's of 7.2 and 7.3, respectively. Each model represents a linear combination of each data stream with different weights. Both calibration models lead to very similar results, with the combination of each data set allowing for a better estimate than either data stream alone. This can been seen in Figure 3 where the estimate for Friday February 3rd from the combined estimates leads to better results than any individual data stream. However through the combination of the data streams, there is a loss in the accuracy in the timing of events. The linear regression with the activity data results in a median arrival time accuracy of 97% however this is diminished through the multiple linear regression to 69% accuracy as shown in Table 1. The departure and peak occupancy times are diminished from 69% to 66% and from 40% to 37%, respectively.

3.2.1.3 Neural Networks

Overall none of the neural network structures, neural networks, recurrent neural networks, LSTM neural networks, consistently outperformed the linear regression models considering a 24 hour training period. The neural networks without memory performed the best yielding overall median MAE's of 12.6 and 10.1 for NN1 and NN3, respectively. The RNN1 recurrent neural network best estimated peak occupancy amongst the neural networks with a median MAE 13.8. The RNN3 model performed the worst of all models with overall and peak MAE's of 21.4, and 21.2 respectively. With respect to timing of events, the NN3 model performed as well as the occupancy estimates based on the univariate regression with the activity data, the only model to do so. The remainder of the models resulted in a degradation in arrival, departure, and peak occupancy accuracies similar to the multiple linear regression and data fusion model structures.

Each of the neural network structure depicts different behavior as shown by the plots of the estimates for a week in Figure 3. The neural networks show similar performance to the linear regressions, keeping the shape of the time series showing behaviors of both the security access data with the long tails and the activity data overestimation on Fridays. The recurrent neural networks have a staggered appearance with the slope of the occupancy estimates often changing from positive to negative from one time step to the next. The LSTM networks result in consistent overestimation and underestimation in the evening hours leading to the worst departure time accuracy of the uncalibrated and calibrated data sets.

3.2.1.4 Variation in Estimates by Day

For the models discussed above, 24 hours of training data were used to estimate occupancy for the remaining 58 days. The analysis showed that the MAE from each training set varied by day. The biggest distinction is between the individual linear regressions and fusion model. To visualise the differences, Figure 4 depicts the MAE resulting from each training data set day for each of the univariate regressions and the data fusion model. The linear regression for the activity data shows some distinct patterns. Higher MAE's were observed for the estimates training with data from December 22, 2017 and December 23rd 2017, the two days before the Christmas holiday. Further higher MAE's are observed for February 3rd 2017, February 10th 2017, February 17th 2017, and March 24th 2017 each of which are Fridays. This indicates that the relationship between computer activity and building occupancy varies by the intensity of activities performed in the building as one would assume less computer activity on these days. The Wifi connectivity and security access univariate regressions indicate days with higher MAE but an overall trend is not clear. The data fusion estimate is a weighted average of each of the estimates, with the largest weight given to the activity

Figure 3: Calibrated Occupancy Counts for each model structure for the week of January 30th 2017. Training Data: Wednesday January 25th 2017
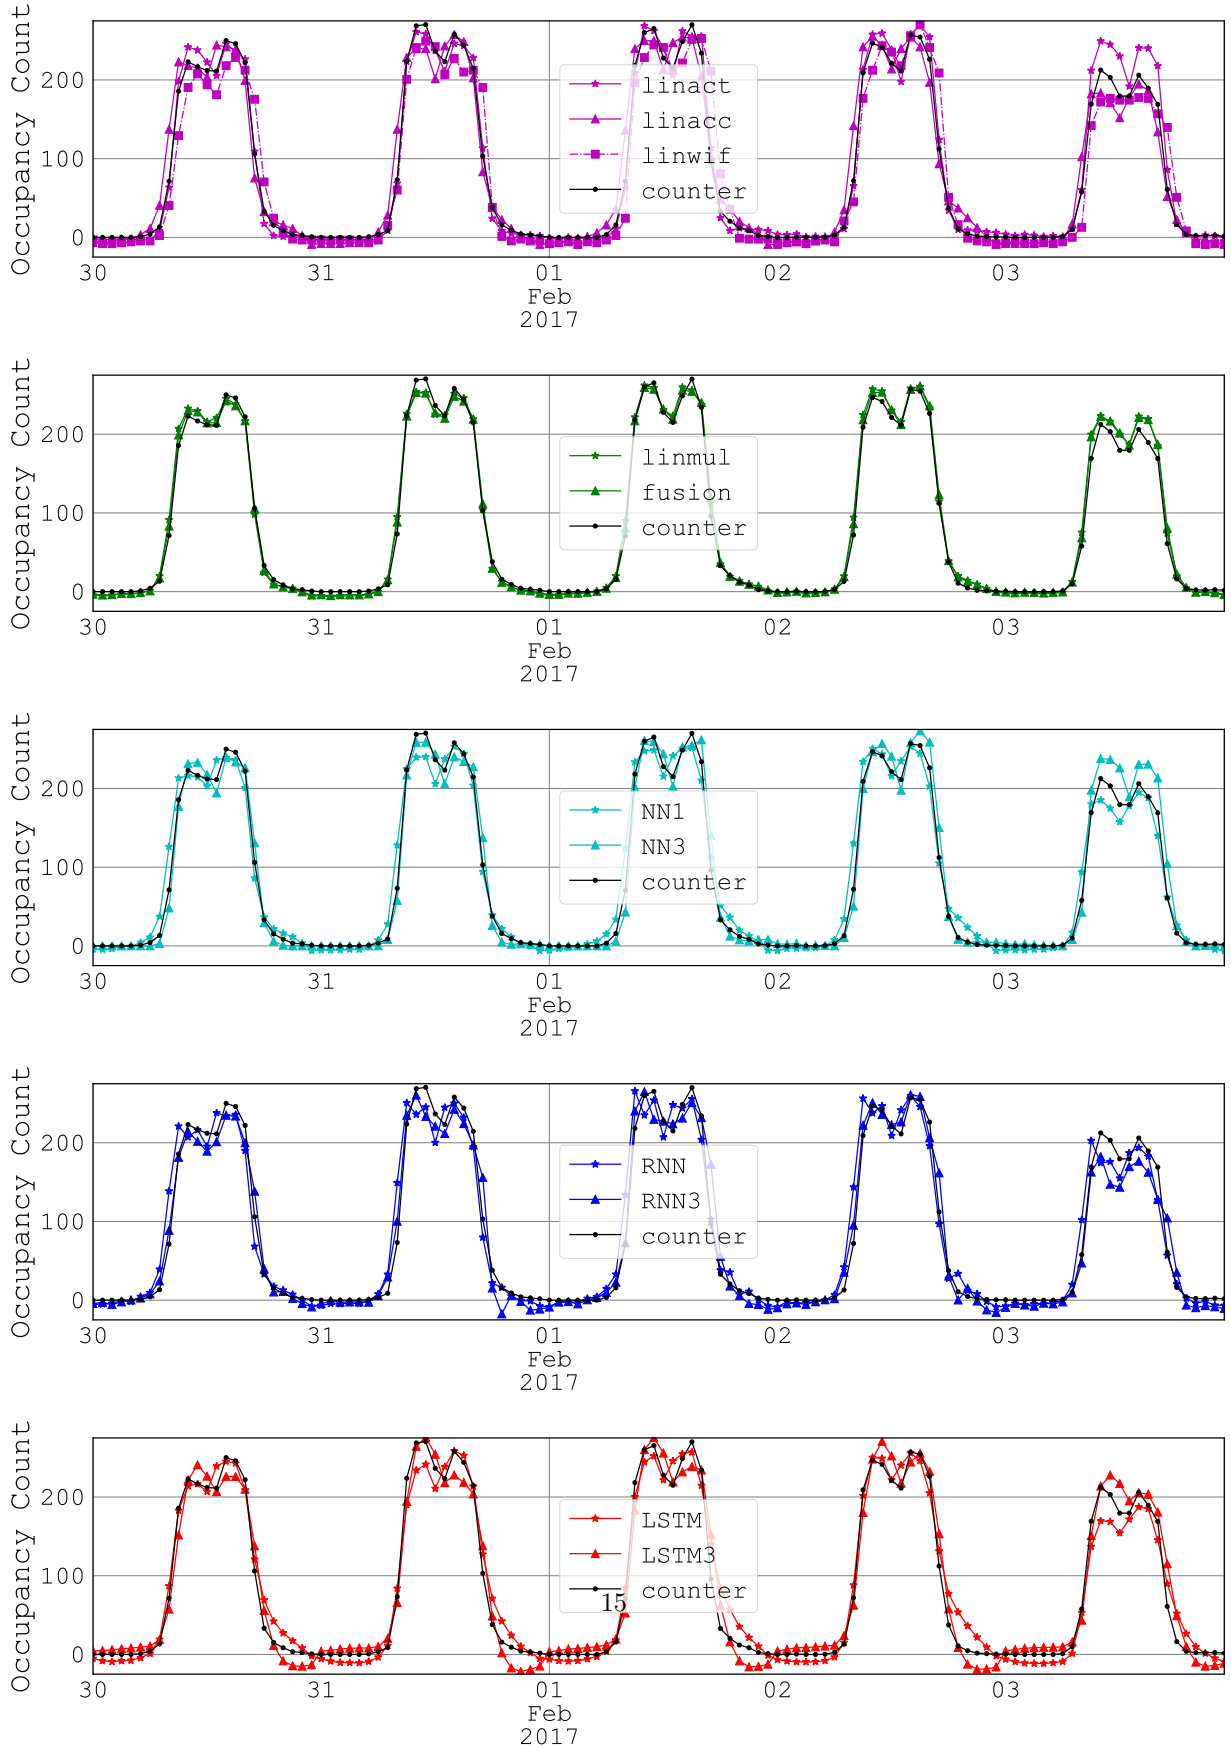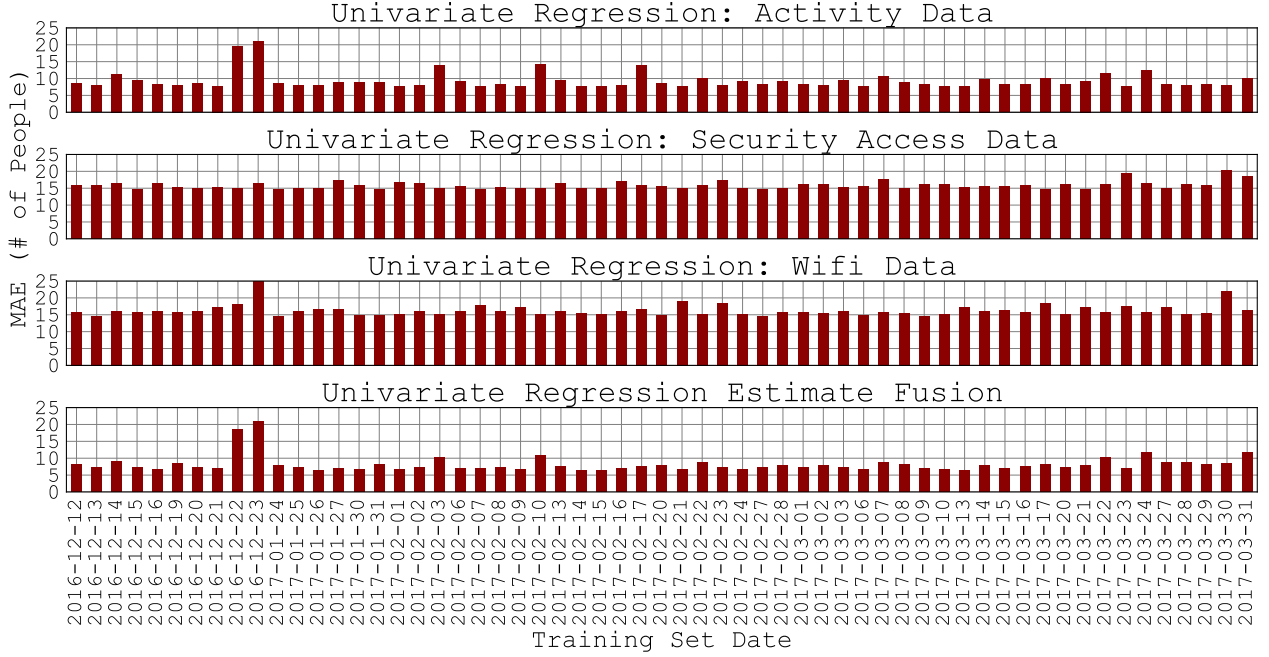
Figure 4: MAE by Training Set Date for each univariate regression and the data fusion model structures

data. Whilst achieving lower MAE's overall, the trend in higher error in the days before Christmas and Fridays persists.

### 3.2.2. Effect of Training Set Length

Performance metrics for each model structure were evaluated for training data sets of 24, 48, and 120 hours. Figure 5 depicts the median values of the performance metrics found for each model structure and training set length. For all model structures, there is a decrease in the median overall and peak occupancy MAE, indicating better performance with longer training sets. In the majority of cases, this reduction is small amounting to a lower MAE of 1 person when increasing the training set length from a day to a week. The cases with larger reductions are those for the neural network structures. The biggest change is for the RNN3 model structure for which the overall MAE reduced from 21.4 to 15.1 from 24 hours to 120 hours, respectively.

The accuracy metrics for the majority of model structures have a small changes as well, however performance increased and decreased for different model structures across the increasing training period. For the arrival time accuracy, with the exception of NN1, the increasing training period reduced the accuracy. The largest change was for the data fusion model whose accuracy reduced from 76% to 69%. For the departure time accuracy, the majority model structures increased accuracy, with the exception of the data fusion model structure, whose departure time accuracy reduced from 69% to 65%. The largest increases in accuracy were found for the NN1, NN3, RNN1, RNN3, and LSTM3 model structures. The median peak occupancy time accuracy for each model structure at 120 hours remained within 2% of the accuracy reported at 24 hour (e.g. 39% at 24 hours and 41% at 120 hours).

4. Discussion

The computer activity data, security access data, and the WiFi connectivity data each contained valuable information on occupancy count. Qualitatively, each data stream is visually similar to the occupancy count from the commercial counter but with a lower magnitude. Without calibration, each data set underestimated

Figure 5: Effect of Time Length of Training Data Set on Performance Metrics
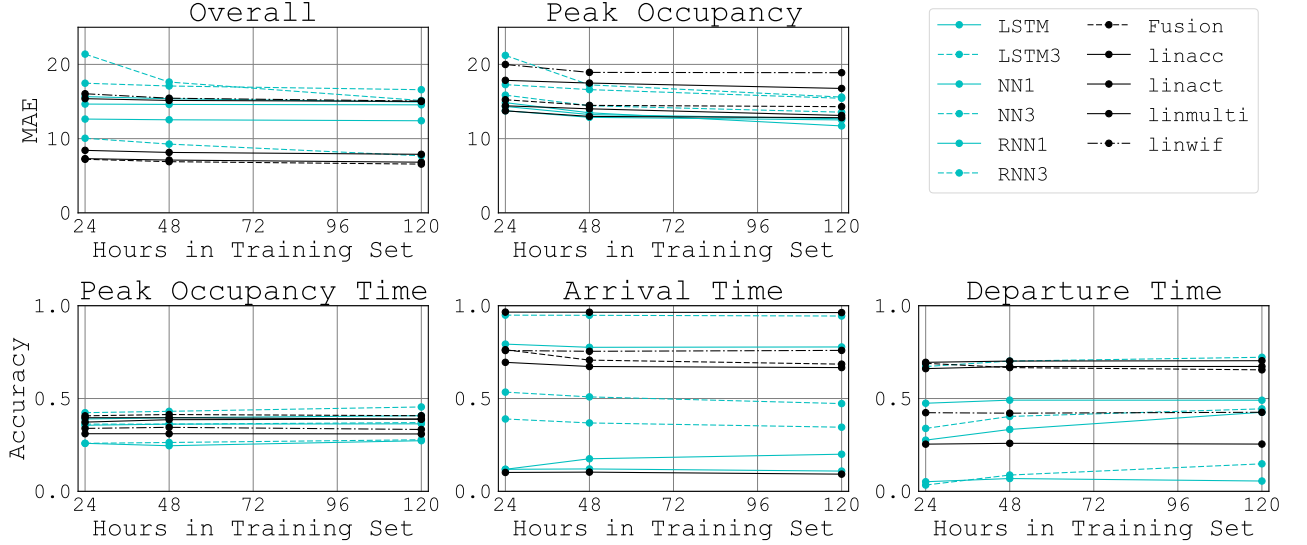
the overall occupancy by 15 to 30 people and during peak times this increased to 30 to 80 people. The activity data , however, was able to estimate the arrival time, i.e. the first time at which more than 17.5 people were in the building, with 97%accuracy. Overall calibration is required to provide more accurate estimates of occupancy count.

With calibration, despite which model, the data sets contained enough information to provide more accurate estimates of occupancy count with respect to the MAE performance indicators, considering 24 hours of training data.

While relatively high arrival time accuracy was achieved, the departure time accuracy for all models was fairly low considering the 1 hour time scale. This is caused by the difference in the morning and evening behavior of the occupants. In the morning, there is a rapid increase in occupancy. The occupancy goes from 10-15 people in 1 hour to 60-70 people in the next. In the evening, the reduction in occupancy is more gradual with occupancy reducing by 10-15 people per hour from 5pm to 7pm. This means that an error of a few people during this time can lead to an inaccurate estimation of departure time. For the peak occupancy timing, a similar explanation can be made. With the bimodal shape of occupancy in office buildings, two peaks arise and, depending on the day, the peak occupancy can be in the afternoon or evening. The MAE for most calibrated model structures was around 15 people which is similar to the difference in magnitude between the two peaks i.e. the morning and afternoon peaks are often times within 15 people of each other. Therefore estimating the time of peak occupancy requires higher accuracy than provided by the ICT data sets.

Increasing the length of the calibration period from 24 hours to 120 hours had only a slight affect on the performance metrics with the exception of the neural network model structures. With a week of training data, the NN3 model was able to increase its performance to be closer to or better than the linear regression with activity data, the data fusion of the linear estimates, and the multiple linear regression. As the neural network models have more coefficients than the linear regressions, it is possible that more data was needed to provide more accurate estimates. This indicates that perhaps with many more data streams available for estimating occupancy, neural networks could yield better performance metrics.

Determining which approach performed best is difficult, as different models performed well in different indicators. Considering 24 hours of training, the two competing models are the data fusion approach and the linear regression with activity data. The data fusion approach performed relatively well across all metrics with the lowest MAE_o. However in comparison to the linear regression with activity data, the arrival time

17

accuracy was diminished. The linear regression with activity performed the best in arrival time accuracy and performed relatively well in the other indicators with the exception of the MAE_p.

Therefore there is a trade off in the approach used. The data fusion requires all three data streams, provides better overall occupancy counts but leads to lower arrival time accuracy. The linear regression with activity data only requires the single data stream, performs well with timing of events but performs worse in estimating peak occupancy count. However, the linear regression with the activity data indicated a large fluctuation in MAE depending on the day used for training, indicating the relationship between computer activity and occupancy could change depending on the intensity of work performed in the building. Using this data stream in isolation could lead to a degradation in accuracy as the work load, work tools or end-use of the organization change.

Overall with these data sets even considering calibration, the ICT data sets were not able to estimate occupancy as well as the commercial counter. The MAE_o and MAE_p for the data fusion and linear regression with activity data are approximately equivalent to an occupancy counter with an overall accuracy of 90%. This opens the question of "how accurate does the occupancy count need to be for improved operation in building heating ventilation and air conditioning systems?".

For occupant driven scheduling of heating ventilation and air conditioning systems, i.e., determining the operating hours from occupancy, error in the estimates could be adjusted for by adding a time to the estimated arrival and departure times. An overestimation of the occupied period would lead to reduced energy savings from the approach and if large enough could increase energy consumption due to extended operating hours. For demand controlled ventilation, a similar adjustment could be made also leading to diminished energy savings from the demand controlled approach. Given systems without demand controlled ventilation are often operated for the maximum occupancy, increasing the energy consumption is unlikely. For model predictive control approaches where occupancy is used to estimate internal loads, inaccurate estimates could lead to decreased thermal comfort as it is possible that the system would not be able to properly compensate for the internal gains. More research is needed to understand the magnitude of the affect of occupancy count inaccuracies on the operation of these systems. Overall the accuracy need from the occupancy counts will be determined by how the data is incorporated into the building operation.

5. Limitations and Areas for Future Work

A key limitation of this work is the context of the case study building. In the analysis, linear regressions, where the data were multiplied by a single factor plus a bias, yielded consistent results. This could be due to the case study setting of an office building with very regular activity patterns. For a building with more variable use, the relationship between ICT system use and occupancy could change over time.

With respect to data collection, very limited analysis was performed on the Wi-Fi data, in that the devices associated with the network were directly taken as occupancy count. It could be that a different indicator of Wi-Fi network activity could yield better performance.

Further data was collected at the 1 hour time scale for the entire building. For maximum savings in building HVAC operations, a smaller time step and higher spatial resolution would be desired. Increasing the spatial resolution of the estimates in the case of Wi-Fi data would require additional analysis and potentially additional wireless access points to determine the position of each device. This would change the classification of the data set from Tier I to potentially Tier III. However the cost to add additional wireless access points may not be viewed as prohibitive.

To draw more general conclusions, future work in the area should attempt to understand the relationship between ICT data sets and occupancy count in buildings with a variety of use cases, with shorter time scales, considering higher spatial resolutions.

6. Conclusions

The aim of this work was to determine if Tier I occupancy sensors could be used to measure occupancy count at scale. Three ICT data sets were evaluated for their performance as compared to a commercial

occupancy counter. With calibration, the ICT data sets were able to closely reflect occupancy count. A 24 hour calibration period is sufficient to achieve accurate measurements, although longer calibration periods allow for more complex models to be used, yielding slightly increased performance. When considering multiple metrics, there are trade offs between the mean absolute error and the timing of events, such as the arrival and departure times. Ultimately, it will be up to the designer of intelligent building control systems to determine which metric is most important and what level of accuracy is needed for system operation.

For the case study building, from the authors perspective, a two layer neural network (NN3) combining the computer activity, WiFi, and security access data sets, using 1 week of training data, resulted in the best performance: MAE overall of 7.7 people, a MAE peak of 15.4 people, arrival time accuracy of 94%, departure time accuracy of 72%, and a peak occupancy time accuracy of 45%. However a linear regression with activity data for only 24 hours yielded similar performance with data fusion and linear regression approaches.

The work described in this paper shows that Tier I implicit occupancy sensors can provide reasonable estimates of occupancy count, although not as accurate as the commercial occupancy counter. This work also highlighted, the potential trade off in accuracy between overall fit and the timing of events. Therefore future work should focus on incorporating and combining several Tier I ICT data sets to enable more accurate spatial and temporal occupancy count estimates in a variety of use cases, whilst considering the performance amongst multiple objectives.

7. Acknowledgments

References

[1] Oswaldo Lucon, Azni Zain Ahmed, Hashem Akbari USA, Paolo Bertoldi, Luisa F Cabeza, Peter Graham, Marilyn Brown, Fonbeyin Henry Abanda, Katarina Korytarova, D Ürge-Vorsatz, A Zain Ahmed, H Akbari, P Bertoldi, L F Cabeza, N Eyre, A Gadgil, L D D Harvey, Y Jiang, E Liphoto, S Mirasgedis, S Murakami, J Parikh, C Pyke, M V Vilariño, R Pichs-Madruga, Y Sokona, E Farahani, S Kadner, K Seyboth, A Adler, I Baum, S Brunner, P Eickemeier, B Kriemann, J Savolainen, S Schlömer, C von Stechow, T Zwickel, and JC Minx. Buildings. Technical report, 2014.

[2] Building control systems CIBSE Guide H. Technical report.

[3] Bill Von Neida, Dorene Manicria, and Allan Tweed. An Analysis of the Energy and Cost Savings Potential of Occupancy Sensors for Commercial Lighting Systems. *Journal of the Illuminating Engineering Society*, 30(2):111–125, jul 2001.

[4] X. Guo, DK Tiller, GP Henze, and CE Waters. The performance of occupancy-based lighting control systems: A review. *Lighting Research & Technology*, 42(4):415–431, dec 2010.

[5] Bart Merema, Muhannad Delwati, Maarten Sourbron, and Hilde Breesch. Demand controlled ventilation (DCV) in school and office buildings: Lessons learnt from case studies. *Energy and Buildings*, 172:349–360, aug 2018.

[6] Amin Mirakhorli and Bing Dong. Occupancy behavior based model predictive control for building indoor climatea critical review. *Energy and Buildings*, 129:499 – 513, 2016.

[7] Miroslava Kavgic, Trent Hilliard, and Lukas Swan. Opportunities for implementation of mpc in commercial buildings. *Energy Procedia*, 78:2148 – 2153, 2015. 6th International Building Physics Conference, IBPC 2015.

[8] Sorin C. Bengea, Anthony D. Kelman, Francesco Borrelli, Russell Taylor, and Satish Narayanan. Implementation of model predictive control for an hvac system in a mid-size commercial building. *HVAC&R Research*, 20(1):121–135, 2014.

[9] Bengea S. Lin Y. Taylor R. Vrabie D. Yuan S. Killough S. Kuruganti T Manges W. Woodworth K. T. Borrelli F. Narayanan, S. and A. Kelman. A wireless platform for energy-efficient building control retrofits. final report estcp project ew-200938, 2012.

[10] Anil Aswani, Neal Master, Jay Taneja, David Culler, and Claire Tomlin. Reducing transient and steady state electricity consumption in hvac using learning-based model-predictive control. *Proceedings of the IEEE*, 100:240–253, 2011.

[11] Bing Dong and Khee Poh Lam. A real-time model predictive control for building heating and cooling systems based on the occupancy behavior pattern detection and local weather forecasting. *Building Simulation*, 7(1):89–106, Feb 2014.

[12] Weiming Shen, Guy Newsham, and Burak Gunay. Leveraging existing occupancy-related data for optimal control of commercial office buildings: A review. *Advanced Engineering Informatics*, 33:230–242, aug 2017.

[13] Bing Dong, Mikkel Baun Kjaergaard, Marilena De Simone, H Burak Gunay, Jakub Dziedzic William O'Brien, Dafni Mora, and Jie Zhao. *Exploring Occupant Behavior in Buildings: Methods and Challenges*. Springer, 2018.

[14] International Energy Agency. *Annex 66 Final Report: Definition and Simulation of Occupant Behavior in Buildings*. The Regents of the University of California (through Lawrence Berkeley National Laboratory) and Tsinghua University, China, 2018.

[15] Tero Kivimäki, Timo Vuorela, Pekka Peltola, and Jukka Vanhala. A Review on Device-Free Passive Indoor Positioning Methods. *International Journal of Smart Home*, 8(1):71–94, 2014.

[16] Thiago Teixeira, Gershon Dublon, and Andreas Savvides. A Survey of Human-Sensing: Methods for Detecting Presence, Count, Location, Track, and Identity. Technical report, Yale University, 2010.

[17] Varick L. Erickson, Miguel Á. Carreira-Perpiñán, and Alberto E. Cerpa. Occupancy Modeling and Prediction for Building Energy Management. *ACM Transactions on Sensor Networks*, 10(3):1–28, may 2014.

[18] Fisayo Caleb Sangoboye and Mikkel Baun Kjærgaard. PLCount: A Probabilistic Fusion Algorithm for Accurately Estimating Occupancy from 3D Camera Counts. In *Proceedings of the 3rd ACM International Conference on Systems for Energy-Efficient Built Environments - BuildSys '16*, pages 147–156, New York, New York, USA, 2016. ACM Press.

[19] M. Berger and A. Armitage. Room occupancy measurement using low-resolution infrared cameras. In *IET Irish Signals and Systems Conference (ISSC 2010)*, pages 249–254. IET, 2010.

[20] Bharathan Balaji, Jian Xu, Anthony Nwokafor, Rajesh Gupta, and Yuvraj Agarwal. Sentinel: occupancy based HVAC actuation using existing WiFi infrastructure within commercial buildings. In *Proceedings of the 11th ACM Conference on Embedded Networked Sensor Systems - SenSys '13*, pages 1–14, New York, New York, USA, 2013. ACM Press.

[21] A. Corna, L. Fontanab, A. A. Nacci, and D. Sciuto. Occupancy Detection via iBeacon on Android Devices for Smart Building Management. In *Design, Automation & Test in Europe Conference & Exhibition (DATE), 2015*, pages 629–632, New Jersey, 2015. IEEE Conference Publications.

[22] Ken Christensen, Ryan Melfi, Bruce Nordman, Ben Rosenblum, and Raul Viera. Using existing network infrastructure to estimate building occupancy and control plugged-in devices in user workspaces. *International Journal of Communication Networks and Distributed Systems*, 12(1):4, 2014.

[23] Ryan Melfi, Ben Rosenblum, Bruce Nordman, and Ken Christensen. Measuring building occupancy using existing network infrastructure. In *2011 International Green Computing Conference and Workshops*, pages 1–8. IEEE, jul 2011.

[24] Davide Calì, Peter Matthes, Kristian Huchtemann, Rita Streblow, and Dirk Müller. CO2 based occupancy detection algorithm: Experimental analysis and validation for office and residential buildings. *Building and Environment*, 86:39–49, apr 2015.

[25] Chaoyang Jiang, Mustafa K. Masood, Yeng Chai Soh, and Hua Li. Indoor occupancy estimation from carbon dioxide concentration. *Energy and Buildings*, 131:132–141, nov 2016.

[26] Tobore Ekwevugbe, Neil Brown, Vijay Pakka, and Denis Fan. Improved occupancy monitoring in non-domestic buildings. *Sustainable Cities and Society*, 30:97–107, apr 2017.

[27] J Hutchins, A Ihler, and P Smyth. Modeling count data from multiple sensors: A building occupancy model. *2nd IEEE International Workshop on Computational Advances in Multi-Sensor Adaptive Processing*, pages 241–244, 2007.

[28] F.C. Sangogboye and M. Kjaergaard. Plcount: A probabilistic fusion algorithm for accurately estimating occupancy from 3d camera counts. *Proceedings of the 3rd ACM International Conference on Systems for Energy-Efficient Built Environments*, pages 147–156, 2016.

[29] J Kuutti, P Sarrikko, and R.E. Sepponen. Real time building zone occupancy detection and activity visualization utilizing a visitor counting sensor network. *11th International Conference on Remote Engineering and Virtual Instrumentation*, pages 219–224, 2014.

[30] Acha S. Shah N. Polak J. Howard, B. Estimating Building Occupancy through ICT Data Streams. In *European Council for an Energy Efficient Economy Summer Study*, 2017.

[31] JM bates and C W J Granger. The combination of forecasts. *Journal of the Operational Research Society*, 20:451–468, 1969.

[32] Allan Timmermann. Chapter 4 Forecast Combinations. In *Handbook of Economic Forecasting*, pages 135–196. 2006.