Deep-learning for Occupancy Detection Using Doppler Radar and Infrared Thermal Array Sensors

M. Abedi[a], and F. Jazizadeh[b]

[a] Charles Edward Via. Department of Civil and Environmental Engineering, Virginia Tech, US
E-mail: jazizade@vt.edu

Abstract –

An accurate, real-time monitoring of the occupancy state of each space is a necessity for applications such as energy-aware smart buildings. In this paper, we have studied the feasibility of using Doppler Radar Sensors (DRS) and Infrared Thermal Array Sensors (ITA) to build an effective occupancy detection framework. The proposed sensor types are cost-effective and protect the privacy of the occupants. We have utilized Deep Neural Networks (DNN) to analyze the sensor data without any need for specialized feature extraction that is necessary for classical machine learning approaches. The results are indicative of the feasibility and the reliability of using both sensor types for detection of the occupancy state. While a threshold-based approach reached an average accuracy of 84.3% and 86% for the DRS and ITA sensors respectively, DNN models were able to achieve average accuracies of 98.9% and 99.96% for the DRS and ITA sensors respectively, thereby demonstrating the feasibility and success of the proposed framework.

Keywords –

Doppler Radar Sensor; Infrared Thermal Array Sensor; Deep Learning; Deep Neural Network

1 Introduction

In the US, residential and commercial buildings are responsible for 39% of the total energy consumption [1]. About 8% of the total electricity consumed in residential and commercial buildings is used for lighting [2]. Heating, Ventilations, and Air Conditioning (HVAC) systems, being responsible for the consumption of 40% of total energy used by the building sector, are also considered to be a major consumer of energy [3]. The combination of these figures underlines the necessity of the notion of energy-aware “smart buildings”.

The nature of the relationship between the demand for HVAC and lighting in the environment (e.g. a room) and the state of occupancy of the environment, has given rise to a growing demand for effective occupancy detection technologies. Several smart HVAC control frameworks rely directly on information regarding the occupancy state of the environment for control purposes [4,5]. Occupancy-based control of the lighting system has also been the subject of research studies [6] and has been widely adopted in practice. The potential for energy saving by enabling occupancy-based control of building systems has been estimated to be as high 25% [7] for HVAC systems and 50% [8] for lighting systems.

A necessary condition for the realization of the true potential of occupancy-based smart control of building operations, is the development of accurate and robust sensing frameworks, easily deployable in a variety of environments. Two of the traditional sensing frameworks for detection of occupancy are such technologies as Passive Infrared sensors (PIR) [9], and image-based technologies [10,11]. However, the PIR sensing framework has been known to suffer from high error rates [12]. PIR sensors require an unobstructed view of the occupant to function and their ability to detect the occupant deteriorated with the increase in distance-to-target [13]. Moreover, the ability of PIR sensors to detect the occupant relies on the existence of a temperature contrast between the occupant and the surrounding environment, thereby resulting in a performance loss in warmer room environments [14]. The mechanism of detection of the occupant for PIR sensors relies on measuring the changes in the temperature contrast of the monitored area, as created by movements of the occupant. Thus, the performance of the PIR sensor is predicated on clear occupant movements, rendering the sensor insensitive to more subtle movements by occupants’ body parts.

An alternative to PIR sensors is the image-based sensing technology. In particular, video-camera-based occupancy detection frameworks are considered to be an effective alternative to PIR sensors. However, implementation of these technologies is accompanied by considerable privacy concerns. As such, researchers have endeavored to propose alternative sensing technologies to address deficiencies of the traditional approaches while respecting occupant privacy.

One type of sensors that we have evaluated in this paper with regards to their applicability for occupancy detection purposes are Doppler Radar Sensors (DRS). DRS sensors rely on motion for detection of the state of the occupancy. DRS sensors can detect the existence of the motion by measuring the change in the frequency of the reflected wave as a result of target motion. However, unlike PIR sensors, DRS sensors are capable of detecting subtle movements of the body such as rotation of the occupants' head, movement of the occupants' hands and arms, and even pulmonary activities, thereby resulting in higher accuracy and reliability.

Another sensor type that has been subject to investigation in the present paper is the Infrared Thermal Array (ITA) sensor. While ITA sensors have been prevalently present in the market, until recently there was a gap in the production resolution of these sensors and the users had to choose between an affordable sensor resolution of up to 64 pixels (8×8 or 4×16) or opt for high resolution sensors at an approximate cost of $200 [15] per sensor. However, in early months of 2018, mid-resolution (32×24) ITA sensors (e.g. model name: MLX90640 [16]) at an affordable cost of under $50 have been introduced to the market. These sensors present a unique opportunity for utilization in an occupancy detection framework as the resolution/cost balance has become reasonable enough to allow for such applications. While our original intention was to build a framework that would utilize both sensors simultaneously to reduce individual error-rates, the high accuracy of each individual sensor convinced us to evaluate them separately.

For the proposed sensing setup to be operationally feasible, the framework must be augmented with an effective data analysis algorithm. By defining a binary state of occupancy for the room, the problem of occupancy detection becomes one of binary classification. Within the multitude of well-established classification algorithms, we have opted to utilize a Deep Neural Network (DNN) model for analysis of sensor data. Utilization of DNN models will allow us to obviate the need for feature-extraction step thereby resulting in an autonomous data analysis framework.

2 Literature Review

Given the growing demand for energy-efficient smart buildings, researchers have endeavored to propose a multitude of occupancy monitoring frameworks consisting of various sensor types and data-analysis strategies.

One of the well-studied sensor types for indoor occupancy monitoring are Passive Infrared (PIR) sensors. PIR sensors measure the infrared light emitted from the object and detect the movements of the source of emission. However, these sensors are not sensitive to subtle and slow movements, which diminishes their capacity to perform as a presence sensor. As such, researchers have endeavored to rectify the aforementioned limitation by means of algorithmic developments and augmentation with other sensor types. For instance Pedersen, et al. [17] have augmented PIR sensors with additional noise, CO_2, Volatile Organic Component (VOC), humidity, and temperate sensors to monitor the state of occupancy of a room. In another study, Dodier, et al. [18] have proposed a belief network approach to analyze the data coming from a network of PIR sensors.

Video cameras have also been used for occupancy detection (mostly occupancy counting) and monitoring. For instance, Hoover and Olsen [10] and Fleuret, et al. [19] have used video cameras to enable tracking of occupied spaces within a room. While implementations of camera-based occupancy detection methods have proven to be accurate [20], utilization of cameras is accompanied by privacy considerations. As such, researchers have endeavored to propose sensing frameworks that are both accurate and compliant with privacy expectations.

Doppler Radar Sensors (DRS) have been studied by researchers as an alternative to traditional occupancy detection frameworks. Like PIR sensors, DRS sensors also detect the motions in their field of view, however, their ability to detect very subtle movements such as those created by pulmonary activities, allows DRS sensors to circumvent some of the important limitations of PIR sensors. Lurz, et al. [21] have demonstrated the feasibility of using DRS sensors for occupancy detection in an experiment that emulated human respiration by means of a linear stage at a distance of 2 m from the sensor. Yavari, et al. [22] have used DRS sensors to detect occupancy by relying on extraction of pulmonary and cardiovascular signatures in the DRS signal while the occupant was either at rest or was moving at different activity levels. One limitation of their study was the constant 1.5 m distance between the radar and the occupant. In the present paper, we have sought to extend the investigation of the capability of DRS sensors to measure occupancy throughout the room without preset conditions such as restricting occupant distance from the sensor to enable room-level service. We have used wide-angle DRS sensors, installed at the ceiling to monitor the state of occupancy in a typical office room.

Another sensor type that has presented a potential for occupancy detection is the Infrared Thermal Array Sensor (ITA). Beltran, et al. [7] utilized an ITA sensor with the resolution of 8×8 to monitor room occupancy state and then used the knowledge of occupancy to more efficiently control the HVAC system operations thereby achieving an annual energy saving rate of 25%. The algorithms utilized in the aforementioned study [7] for the interpretation of the ITA sensor outputs are K-Nearest Neighbor (KNN), Linear Regression, and an Artificial Neural Network. In another study, Tyndall, et al. [23] used an ITA sensor with the resolution of 4×16 to estimate the state of occupancy inside the room. The algorithms used for analysis of data consist of a number of standard classification algorithms such as Support Vector Machine, and KNN [23].

Recent industry advancements have given rise to semi-high resolution ITA sensors at an affordable price of approximately $50 (e.g. MLX90640 [16] with a resolution of 32×24). This development could result in an increased potential for utilization of ITA sensors for non-intrusive occupancy-monitoring applications. As such, in this paper we have evaluated the performance of the MLX90640 [16] sensor for real-time monitoring of the occupancy state in a typical office room environment. Additionally, by implementing a Deep-Learning solution for the analysis of sensor data we have eliminated the feature extraction step, thereby rendering the data analysis step autonomous.

3 Methodology

Presented in this section is the discussion of two aspects of the occupancy detection framework. In the first part we have discussed the overview of the sensing framework. In the subsequent section, the data analysis and occupancy inference approach have been discussed.

3.1 Sensing System Setup and Data Acquisition

As noted, in this study, we have investigated the potential of two distinct sensor types for occupancy detection applications. The first sensor type is the Doppler Radar Sensor (DRS). We have used the RFBeam K-LC3 [24] wide angle doppler transceiver. With dimensions of 25×25 mm and a weight of 5 g the K-LC3 DRS sensor is a cost-effective and practical solution for occupancy detection. Another desirable feature of the selected sensor model is the wide field of view (132º×138º), which allows for coverage of larger areas, compared to other sensor models with narrower fields of view. In Figure 1, the coverage area for the DRS sensor (as well as the ITA sensor) that has been installed in a typical US office room (7.5×4 m) is presented. The sensor is assumed to have been installed on the ceiling with the height of 2.75m (9ft) which is typical in the US. The coverage area has been calculated such that it at least captures the parts of the occupant's body at or under the desk level (0.8 m from the floor), i.e. occupant's hands, lower torso, and legs. As seen in Figure 2., we have installed the DRS sensor on the ceiling close to the air-vent. The DRS sensor has been connected to a SR 560 preamplifier [25] to amplify the output signal by a gain factor of 5×10^4. The amplifier also filters out the frequency content above 1kHz since the frequency content above that threshold is highly unlikely to have been generated due to the motions of the human subject. This sampling rate also helps avoid aliasing error due to possible presence of high-frequency signals associated with the existing equipment in the room (e.g. computer fans and HVAC).

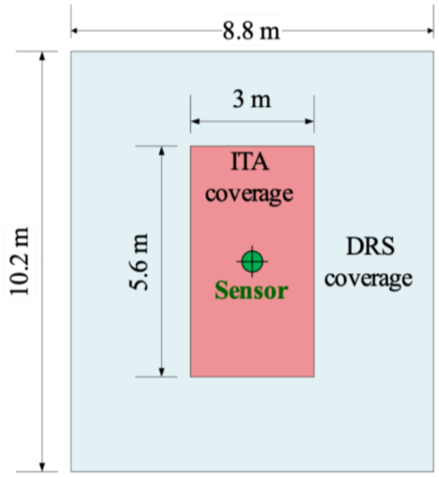

Figure 1. Space covered by each sensor type

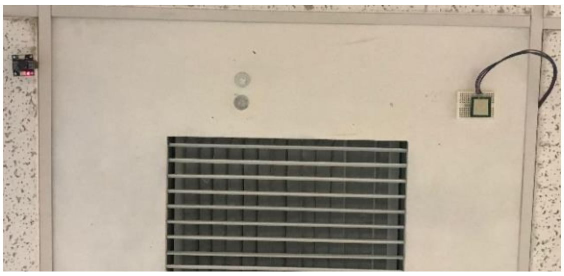

Figure 2. Sensors installed (ITA sensor on the left and DRS sensor on the right)

For data acquisition, we used a National Instrument USB-6001 Multifunction I/O Device [26] to record the amplified signal at a sampling rate of 1kHz. The signal was processed through a zero-phase notch filter to remove the frequency harmonic contents, associated with powerline-noise (i.e. 60 Hz, 120 Hz, …). These frequencies in the signal are usually generated due to the existence of AC powerlines (60 Hz) in the vicinity of the setup.

The second sensor type is the MLX90640 [16] Infrared Thermal Array (ITA) sensor. This ITA sensor has a field of view of 110º×75º divided into 32×24 pixels respectively, where each pixel reads the average temperature of the objects, covered in the view-window of that pixel. We have used a Raspberry Pi 3 [27] computer for data acquisition from the ITA sensor through the I2C serial protocol. We have installed the ITA sensor on the ceiling, close to the air vent as can be seen in Figure 2. The transfer of sensor readings to the on-site computer is done through wire-less connection.

3.2 Occupancy Inference Framework

Each data-point, acquired through our experiments, consists of a 10-second DRS signal @1kHz and a single ITA sensor temperature reading. The DRS data is represented as a time-series and the ITA sensor output is a 2D thermal image. Additionally, each data-point has a binary label (0,1) that determines the state of the occupancy during the 10 second interval associated with each data-point. Through this description, the task of occupancy detection has become one of binary classification (i.e. unoccupied, occupied).

In this study, we have investigated two different occupancy inference methods for the classification problem. In pursuit of an unsupervised method of inference, similar to methods commonly used for PIR sensors, the first approach is a threshold-based decision tree method, that has been used to establish a base-line. Thus, if at any point in time, the DRS time-series exceeds a certain threshold (regardless of the signal amplitude sign), the model will output a positive occupancy state for the room and vice versa. Similarly, if the temperature reading at any of the ITA sensor's output pixels exceeds a certain threshold, the model will output a positive occupancy state for the room and vice versa. The threshold has been selected by using a decision tree of depth 1. The implementation of the decision tree was performed by utilizing the Scikit-Learn [28] library in python 3.65.

As the second class of inference method, we have proposed and evaluated a Deep Neural Network (DNN) solution for the task of occupancy inference based on sensor data. Utilization of a DNN model brings about a number of advantages. Firstly, by utilizing the Deep Learning model, we obviate the need for specialized feature extraction, because the task of feature extraction in a DNN model is performed automatically by the initial layers. Moreover, combining information of fundamentally different nature (in our case a DRS time series and an ITA 2D temperature array) is a relatively effortless task to achieve using DNN models, and it does not require further feature engineering. The structure of the DNN, used for analysis of the DRS time series is as presented in Table 1. Given the high number of trainable parameters in DNNs, there is often a need for a large training data-set to facilitate training and avoid overfitting. However, in some real-world applications such as ours, the size of the data-set is limited. As such, we have opted to use two Dropout layers with a dropout rate of 50% to help avoid overfitting of the model to the training data. The loss function used for training of the model is the 'binary cross-entropy' measure, which is a common choice for binary classification problems.

Table 1. DNN structure for analysis of DRS data with 223,481 trainable parameters.

Layer Type	Layer Size / Drop Rate	Filter Size	Activation Function
Conv1D	128	7	Relu
MaxPool1D	-	2	-
Conv1D	64	6	Relu
MaxPool1D	-	2	-
Conv1D	32	5	Relu
MaxPool1D	-	2	-
Conv1D	16	4	Relu
MaxPool1D	-	2	-
Conv1D	8	4	Relu
MaxPool1D	-	2	-
Flatten	-	-	-
Dropout	50%	-	-
Dense	64	-	Relu
Dropout	50%	-	-
Dense	32	-	Relu
Dense	1	-	Sigmoid

As shown in Table 1, the initial layers consist of relatively large 1D convolutional layers with larger filter sizes. The concatenation of convolutional layers with a shrinking size through the depth results in a phenomenon, through which the first layers are trained to perform generic feature extraction with lesser relevance to the data labels and more relevance to the input data itself. Conversely, the subsequent layers have more and more relevance to the label information. As such, after training the model, one could freeze the initial feature extraction layers, and reuse them for other models, thereby reducing the computational cost of training the new models. This is another helpful feature of DNNs that allows multiple models to share a common body of knowledge. This is an important feature of the model toward generalizability. In occupancy detection, unsupervised models with high accuracy and reliability are preferred.

In Table 2, the structure of the DNN used for analysis of ITA sensor output has been presented. Similar to the previous model, we have utilized Dropout layers to avoid overfitting to the training data. The initial layer of this model can also be used for transfer-learning purposes to obviate the need to re-train feature-extraction layers in a new DNN model. Similarly, the loss function has been chosen to be 'binary cross-entropy'.

Table 2. DNN structure for analysis of ITA data with 70,433 trainable parameters

Layer Type	Layer Size / Drop Rate	Filter Size	Activation Function
Conv2D	64	(3,3)	Relu
MaxPool2D	-	(2,2)	-
Conv2D	32	(3,3)	Relu
MaxPool2D	-	(2,2)	-
Flatten	-	-	-
Dropout	50%	-	-
Dense	64	-	Relu
Dropout	50%	-	-
Dense	32	-	Relu
Dense	1	-	Sigmoid

Both models have been implemented by utilizing the Keras [29] library with the TensorFlow [30] backend. Before feeding the data to the model, we have subtracted the mean from the ITA sensor readings. This has been done in order to stabilize the model and also to break possible time dependent relationships among data-points, since data-points that have been recorded close to each other could have a similar mean temperature reading. Thus, the ITA readings become temperature differentials rather than actual temperature readings. An example of these readings has been presented in Figure 3.

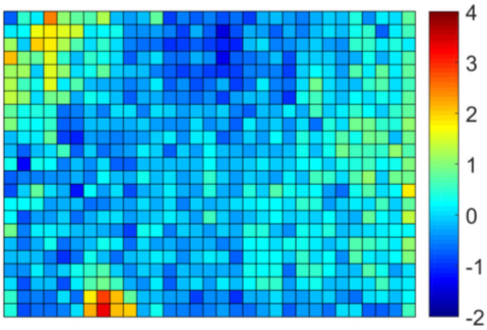

(a) The room is occupied

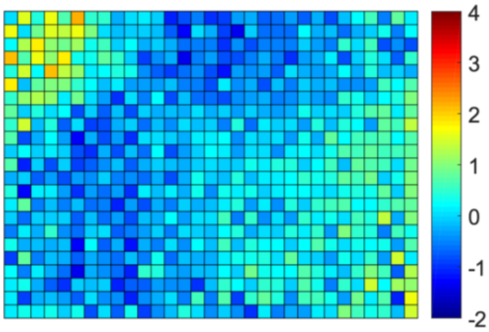

(b) The room is unoccupied

Figure 3. Temperature differentials (°C)

As can be seen in Figure 3, in both the occupied and unoccupied case, the upper left corner is warmer than the average. This is due to the presence of a personal computer at the corresponding location, which serves as a thermal noise. In Figure 4 an example of DRS sensor readings has been presented.

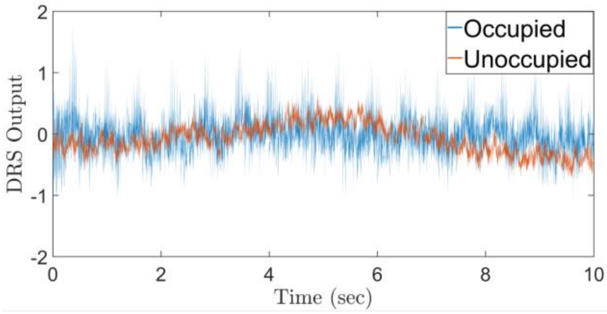

Figure 4. DRS sensor reading

Given that DRS sensor readings for both occupied and unoccupied cases were mostly in the [-1,1] range, no normalization was deemed necessary.

4 Results

In order to train the model to detect the state of occupancy in the environment, we need to provide the model with a data-set that includes both occupied and unoccupied states. The data associated with an unoccupied state, have been collected during the weekend. As for data associated with a positive state of occupancy, the human subject has sat at two different locations within the field of view of the ITA sensor.

The acquired dataset for this study consists of 1000 data-points, where each data-point contains an ITA sensor reading and 10 seconds of DRS signal (but only one sensor's reading will be given to each model, i.e. either DRS or ITA). In 500 of these cases, the room was unoccupied and in the other 500 cases the room was occupied. The evaluation of the models was performed through a 3-fold cross-validation. As noted, a simplified, threshold-based model was initially used to establish a baseline against which the performance of the DNN model is to be compared. In training of all models, the loss function has penalized equally against both false-positive and false-negative error types. In future studies, the loss function should be modified so as to take into account the relative importance of the two error types for the intended application.

The threshold-based model used to analyze the DRS signal has achieved an average accuracy of 84.3%. By contrast, the DNN model used for the analysis of the DRS signal has achieved an average accuracy of 98.9%. The training process of the model consisted of 10 epochs (batch size of 10) with a total elapsed time of 8 minutes on an Intel Xeon E5-1620 V4 CPU [31].

The threshold-based model for the analysis of the ITA sensor resulted in an average accuracy of 86%. By contrast, the DNN model has been able to reach an average accuracy of 99.96%. Training of DNN model used for analysis of ITA sensors consisted of 15 epochs (batch size of 10) with a total elapsed time of 6 seconds on an Intel Xeon E5-1620 V4 CPU [31].

In order to investigate the performance of the DNN model used for analysis of ITA sensor data, we retrieved the activation values of each layer (on the trained model) when the data-point associated with Figure 3.(a) was passed to the model. The filters in a convolutional layer seek to learn a pattern within the input data, and the activation values associated with each filter layer represents the level of a match that exists between a given input data and the pattern that the particular filter has learned. If the pattern of interest for that particular filter exists in certain areas of the data, the activation values for those areas will be higher. In Figure 5, we have presented the activation outputs for two of the filters in the third layer of the DNN model used for analysis of ITA data, i.e. the Conv2D layer with the size of 32.

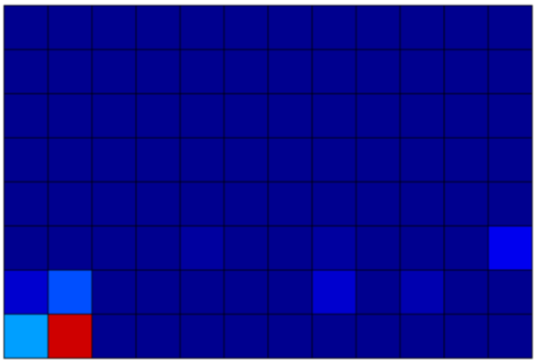

(a)

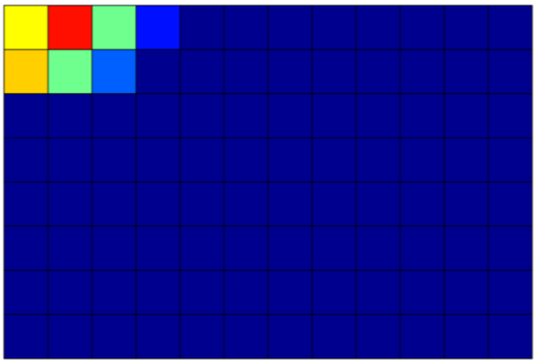

(b)

Figure 5. Activation outputs from two filters in the third layer of the DNN model (i.e Conv2D with the layer size of 32)

By comparing Figure 5.(a) with Figure 3.(a), it becomes apparent that the filter associated with these activation values has been trained to exclusively extract the thermal signatures of a human occupant since the activation values are high at the area associated with the location of the occupant. Interestingly, comparison of Figure 5.(b) with Figure 3.(a) reveals that the filter associated with these activation values has been trained to exclusively learn the thermal signature of the environmental noises (in this case the heat emissions from a personal computer in the environment). The mechanism of learning to detect both the noise and the actual occupant-related signatures in the data could be responsible for the high accuracy of the DNN model.

5 Limitations and Future Work

One limitation of this study was the limited size of the train and test data-sets. Moreover, the results are limited to experiments under a single environmental setting (i.e. one room and one occupant). In our future research we will evaluate the potential of the proposed framework under a multitude of experimental settings. Moreover, we will be investigating the potential of the proposed framework to count the number of occupants within each room.

Another line of research which will be the subject of our future studies is the potential to jointly learn from multiple sensor types to achieve synergistic combinations. For instance, in a room with occasionally high thermal noise, the DRS data could help reduce the possible errors in the performance of the ITA-based system. Similarly, if the degree of non-occupant related activities increases to a level that it would affect the performance of the DRS sensor, the information from ITA sensor could compensate for the possible increase in the DRS-based occupancy detection system's error rate.

6 Conclusion

A reliable and accurate occupancy detection framework is in high demand for realization of such applications as smart buildings. In this paper, we have proposed a framework for the detection of indoor occupancy based on cost-effective Doppler Radar Sensors (DRS) and cost-effective high-resolution Infrared Thermal Array Sensors (ITA). We have presented a Deep Neural Network solution for the analysis of the sensor data. We have further evaluated the performance of the framework by conducting real-work experiments in a typical office room. In order to establish a base-line for the performance of the proposed framework, we utilized a simplified threshold-based approach. The threshold-based approach achieved an average accuracy of 84.3% and 86% for the DRS and ITA sensors respectively. The proposed DNN framework has demonstrated an accuracy of 98.9% via the DRS sensor and 99.96% via the ITA sensor.

Acknowledgement

This material is based upon work partially supported by the National Science Foundation under grant #1663513. Any opinions, findings, and conclusions or recommendations expressed in this material are those of the authors and do not necessarily reflect the views of the National Science Foundation.

References

[1] U.S. Energy Information Administration. How much energy is consumed in U.S. residential and commercial buildings? On-line: https://www.eia.gov/tools/faqs/faq.php?id=86&t=1 Accessed: 24/01/2019

[2] U.S. Energy Information Administration. How much electricity is used for lighting in the United States? On-line: https://www.eia.gov/tools/faqs/faq.php?id=99&t=3 Accessed: 24/01/2019

[3] Department of the Environment and Energy of the Australian Government. HVAC factsheet - Energy breakdown On-line: https://www.energy.gov.au/publications/hvac-factsheet-energy-breakdown Accessed: 24/01/2019

[4] Agarwal Y. and Balaji B. and Gupta R. and Lyles J. and Wei M., Weng T. Occupancy-driven energy management for smart building automation. In *Proceedings of the 2nd ACM Workshop on Embedded Sensing Systems for Energy-Efficiency in Building*, pages 1-6, Zurich, Switzerland, 2010,

[5] Nikdel L. and Janoyan K. and Bird S.D., Powers S.E. Multiple perspectives of the value of occupancy-based HVAC control systems. *Building and Environment*, 129:15-25, 2018

[6] Zou H. and Zhou Y. and Jiang H. and Chien S.-C. and Xie L., Spanos C.J. WinLight: A WiFi-based occupancy-driven lighting control system for smart building. *Energy and Buildings*, 158:924-938, 2018

[7] Beltran A. and Erickson V.L., Cerpa A.E. ThermoSense: Occupancy Thermal Based Sensing for HVAC Control. In *Proceedings of the 5th ACM Workshop on Embedded Systems For Energy-Efficient Buildings*, pages 1-8, Roma, Italy, 2013,

[8] Dubois M.-C., Blomsterberg Å. Energy saving potential and strategies for electric lighting in future North European, low energy office buildings: A literature review. *Energy and Buildings*, 43 (10):2572-2582, 2011

[9] Kaushik A.R., celler B.G. Characterization of Passive Infrared Sensors for Monitoring Occupancy Pattern. In *2006 International Conference of the IEEE Engineering in Medicine and Biology Society*, pages 5257-5260, 2006,

[10] Hoover A., Olsen B.D. A real-time occupancy map from multiple video streams. In *Proceedings 1999 IEEE International Conference on Robotics and Automation (Cat. No.99CH36288C)*, pages 2261-2266 vol.2263, 1999,

[11] Bhatia B. and Oates T. and Xiao Y., Hu P. Real-time identification of operating room state from video. In *Proceedings of the 19th national conference on Innovative applications of artificial intelligence - Volume 2*, pages 1761-1766, Vancouver, British Columbia, Canada, 2007,

[12] Srinivasan S. and Pandharipande A., Caicedo D. Presence detection using wideband audio-ultrasound sensor. *Electronics Letters*, 48 (25):1577-1578, 2012

[13] Hodges L. Ultrasonic and Passive Infrared Sensor integration for dual technology user detection sensors.

[14] Beckwith D., Hunter-Zaworski K. Passive pedestrian detection at unsignalized crossings. *Transportation Research Record: Journal of the Transportation Research Board*, (1636):96-103, 1998

[15] FLIR Systems Inc. Lepton LWIR Micro Thermal Camera Module On-line: https://www.flir.com/products/lepton/?model=500-0763-01 Accessed: 24/01/2019

[16] Melexis. Far infrared thermal sensor array (32x24 RES) On-line: https://www.melexis.com/en/product/MLX90640/Far-Infrared-Thermal-Sensor-Array Accessed: 24/01/2019

[17] Pedersen T.H. and Nielsen K.U., Petersen S. Method for room occupancy detection based on trajectory of indoor climate sensor data. *Building and Environment*, 115:147-156, 2017

[18] Dodier R.H. and Henze G.P. and Tiller D.K., Guo X. Building occupancy detection through sensor belief networks. *Energy and Buildings*, 38 (9):1033-1043, 2006

[19] Fleuret F. and Berclaz J. and Lengagne R., Fua P. Multicamera People Tracking with a Probabilistic Occupancy Map. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 30 (2):267-282, 2008

[20] Tomastik R. and Narayanan S. and Banaszuk A., Meyn S. Model-Based Real-Time Estimation

of Building Occupancy During Emergency Egress. pages 215-224, Berlin, Heidelberg, 2010,

[21] Lurz F. and Mann S. and Linz S. and Lindner S. and Barbon F. and Weigel R., Koelpin A. A low power 24 GHz radar system for occupancy monitoring. In *2015 IEEE Radio and Wireless Symposium (RWS)*, pages 111-113, 2015,

[22] Yavari E. and Jou H. and Lubecke V., Boric-Lubecke O. Doppler radar sensor for occupancy monitoring. In *2013 IEEE 13th Topical Meeting on Silicon Monolithic Integrated Circuits in RF Systems*, pages 216-218, 2013,

[23] Tyndall A. and Cardell-Oliver R., Keating A. Occupancy Estimation Using a Low-Pixel Count Thermal Imager. *IEEE Sensors Journal*, 16 (10):3784-3791, 2016

[24] RFbeam Microwave GmbH. K-LC3 Wide Angle Doppler Transceiver On-line: https://www.rfbeam.ch/product?id=6 Accessed: 24/01/2019

[25] Stanford Research Systems. Low Noise Voltage Preamplifier On-line: https://www.thinksrs.com/products/sr560.html Accessed: 24/01/2019

[26] National Instruments. USB-6001 Multifunction I/O Device On-line: http://www.ni.com/en-us/support/model.usb-6001.html Accessed: 24/01/2019

[27] The Raspberry Pi Foundation. Raspberry Pi 3 On-line: https://www.raspberrypi.org/products/raspberry-pi-3-model-b-plus Accessed: 24/01/2019

[28] Pedregosa F. and Varoquaux G. and Gramfort A. and Michel V. and Thirion B. and Grisel O. and Blondel M. and Prettenhofer P. and Weiss R., Dubourg V. Scikit-learn: Machine learning in Python. *Journal of machine learning research*, 12 (Oct):2825-2830, 2011

[29] Chollet F. Keras: The python deep learning library. *Astrophysics Source Code Library*, 2018

[30] Abadi M. and Barham P. and Chen J. and Chen Z. and Davis A. and Dean J. and Devin M. and Ghemawat S. and Irving G., Isard M. Tensorflow: a system for large-scale machine learning. pages,

[31] Intel Corporation. Intel® Xeon® Processor E5-1620 v4 On-line: https://ark.intel.com/products/92991/Intel-Xeon-Processor-E5-1620-v4-10M-Cache-3-50-GHz- Accessed: 24/01/2019