Anomalous occupancy sensor behavior detection in connected indoor lighting systems

Giulia Violatto[1,2], Ashish Pandharipande[1], Shuai Li[1], Luca Schenato[2]
[1]Signify, High Tech Campus, Eindhoven, The Netherlands
[2]Department of Information Engineering, University of Padova, Italy
Email: {giulia.violatto, ashish.p, shuai_li}@signify.com

***Abstract*—We consider the problem of classifying anomalous occupancy sensor behavior in connected indoor lighting systems. Anomalous occupancy sensor behavior may occur in the form of either a high number of false alarms (type-1 anomalies) or missed detection (type-2 anomalies). We consider a supervised machine learning approach to determine whether the detection signal of an occupancy sensor is normal, or exhibits type-1 or type-2 anomalies. We devise occupancy signal features in the time and frequency domains and employ a random forest classifier to perform 3-class classification. The proposed method is evaluated using motion sensor data from an office building, and is shown to have higher true positive rate and a lower false positive rate in comparison to an unsupervised k-means method and a random forest classifier with a single signal energy feature.**

***Index Terms*—Occupancy sensors, Connected lighting, Random forest classifier.**

I. Introduction

There is increased adoption of connected lighting systems for indoor lighting applications. These systems have sensing, control and connectivity functions integrated into luminaires. Conventionally, these functions have been used for smart lighting control to reduce energy consumption and enhance visual comfort of occupants in a building [1], [2], [3], [4]. Smart lighting control system use occupancy sensors and light sensors respectively to adapt the amount of artificial light output to changing occupancy and daylight levels. With granular networked sensing, e.g. with sensors in each luminaire of a connected lighting system, energy reduction up to 80% was shown in office buildings. Occupancy sensor data may also be used for actuating other building systems such as heating, ventilation, and air conditioning (HVAC). As such, proper functioning of occupancy sensors is critical in operations of a number of building control systems.

In this paper, we consider the problem of detecting anomalies in occupancy sensor behavior. A first anomaly type may be in the form of large number of false alarms. This can result due to misconfigurations in detection threshold (being too low for the environment), software issues, bad sensor placement (field-of-view out of the room into a corridor) etc. A second anomaly type may be in the form of large number of missed detection by an occupancy sensor in a room, while a user is actually present. This can occur due to misconfigurations in detection threshold (being too high for the environment), software issues, bad sensor placement (field-of-view obstructed or away from user activity) etc. In the context of lighting, the first type of anomaly may lead to luminaires turning on even when a room is unoccupied leading to lighting energy wastage, while the second type of anomaly leads to luminaires being unresponsive to user presence with the possible consequence of user dissatisfaction with lighting controls.

We consider a supervised machine learning approach to detect anomalies in occupancy sensor signals and pose it as a 3-class classification problem. Towards this end, we consider different features in the time and frequency domain to characterize a normally functional occupancy sensor and a sensor with behavioral anomalies. In a training phase, a random forest classifier [5] is trained to determine informative signal features. This classifier is then used during regular operation to classify sensor behavior as normal, having type-1 or type-2 anomalies.

To the best of our knowledge, the problem of detecting occupancy sensor signal anomalies has not been considered earlier. Most works in lighting literature have considered the use of occupancy sensor data for lighting controls. Works [2], [3], [6] have considered the use of occupancy data for lighting controls to adapt dimming levels of luminaires according to granular occupancy information. Works have also considered occupancy data for improved building energy management and in applications beyond lighting like space management. The use of occupancy data has been explored in applications like space management and energy prediction [7], [8]. The problem of occupancy modeling has been considered in [9], [10], [11], for simulating energy savings in HVAC in commercial buildings. The value of sensor and lighting data to realize new applications and services in connected building eco-systems has been explored in [7], [12], [13], [14]. The problem of detecting commissioning changes in lighting using occupancy sensor data was addressed in [15].

II. System description and problem scope

We consider a connected lighting system installed in an office building. The building has multiple rooms, where each room may have one or more lighting control zones. A lighting control zone consists of one or more luminaires that contribute to illumination in the space. Some or all of these luminaires may be smart luminaires, with co-located occupancy and light sensors that respectively determine user presence states and ambient light levels.

A typical occupancy sensor used in lighting control systems is a Passive InfraRed (PIR) motion sensor. Such a sensor reports an occupied state if motion is detected in its field-of-view (FoV); otherwise, it reports an unoccupied state.

Binary occupancy signals from three sensors are shown in Fig. 1 over four business days. In this example, binary occupancy detection is shown every minute. Sensor 1 has normal behavior, with different detection behaviors over the business days reflecting different occupancy patterns over time. Sensor 2 exhibits type-1 anomalies with an unusually high number of occupancy detection, due to false alarms. Sensor 3 exhibits type-2 anomalies with a very low number of occupancy detection, due to many missed detection.

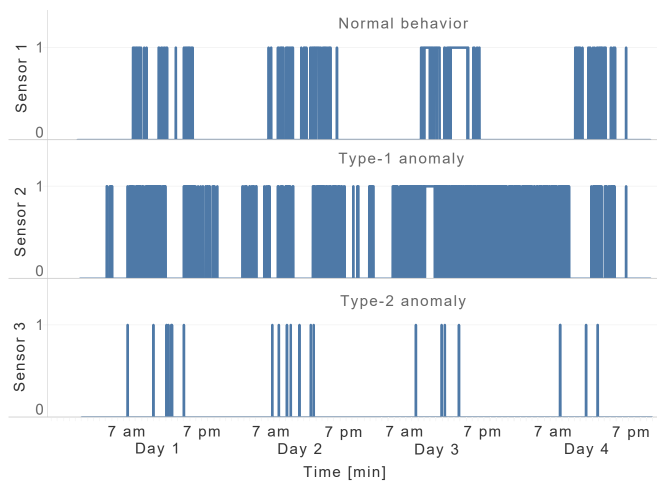

Fig. 1: Occupancy sensor data from three PIR sensors over one week.

Occupancy sensor data with time stamps and sensor IDs are collected at a backend for further analysis. Along with this, commissioning mapping in the form of the mapping of sensor IDs to room IDs and room types (e.g. cell office, meeting room) is also collected and available at the backend.

Consider that the classification task is done using a signal record of N samples of occupancy detection from a sensor. Denote a sensor signal record $\{x_i\}$ to be

$$\{x_i\} = \{x_i[0], x_i[1], \ldots, x_i[N-1]\}.$$

Here i denotes the record index and $x_i[n] \in \{0, 1\}$ denotes the occupancy detection in sensor signal record $\{x_i\}$ at time instant n. For each signal record, the classifier needs to decide whether the behavior is normal, or corresponds to a type-1 or type-2 anomaly. We shall assume that in an initial period after system deployment, sensors have been labeled by a human expert to belong to one of the three classes.

III. Proposed Method

We propose a machine learning approach to detect anomalies in occupancy sensor behavior, which consists of the following steps:

[A] design discriminative time and frequency domain signal features to aid in classification of occupancy sensor behavior;

[B] using the defined features, we tune a random forest classifier to perform 3-class classification between normal behavior, type-1 and type-2 anomalies.

A. Signal features

We consider the following signal features in the time and frequency domain.

1) Time domain features:

i. Signal energy in business and non-business period: Denote the window of time corresponding to the business period by $\mathscr{B}$ and the non-business period by $\mathscr{B}'$. Then the signal energy of sensor signal record $\{x_i\}$ over the two periods is given by

$$E_i^b = \sum_{n \in \mathscr{B}} x_i[n], \quad (1)$$

$$E_i^{b'} = \sum_{n \in \mathscr{B}'} x_i[n]. \quad (2)$$

We plot $E_i^{b'}$ for different signal records, each record taken over a day, in Fig. 2. Signal records with type-1 anomaly generally show high values, although signal records with normal behavior also have some signal energy in instances where users may be working earlier or later into the day.

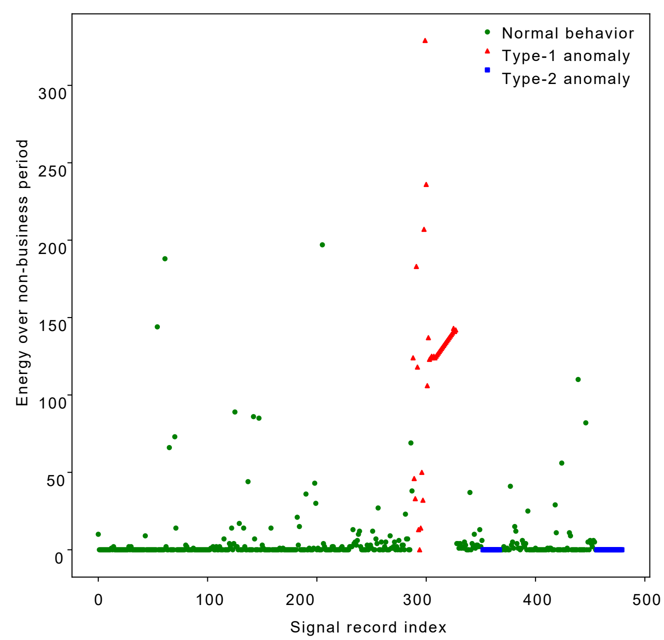

Fig. 2: Energy over non-business period.

ii. Unoccupancy duration: The distribution of the length of unoccupied slots can reveal whether a sensor is normal, has type-1, or type-2 anomaly. Specifically, for signal record $\{x_i\}$, the number of unoccupancy slots of length L is given by

$$u_i^L = \sum_{n,m} \lfloor \frac{m-n-1}{L} \rfloor : x_i[n] = 1 \text{ or } n = 0,$$
$$\text{and } x_i[m] = 1 \text{ or } m = N-1,$$
$$\text{and } x_i[q] = 0 \text{ for } n < q < m, \quad (3)$$

for $0 \leq n < m \leq N-1$. Intuitively, we would expect a sensor functioning normally to have unoccupied slots of a certain duration characterized by occupancy patterns over the business period. A sensor with type-1 anomaly would have too few slots of unoccupancy over long time periods, while a sensor with type-2 anomaly may be expected to have a large number of slots of unoccupancy of a specific length. As example, consider the feature u_i^L with unoccupancy duration of length 60 min shown in Fig. 3 for the signal records. Signal records with type-1 anomaly tend to have lower number of slots and signal records with type-2 anomaly have higher number of such slots as compared to signal records with normal behavior.

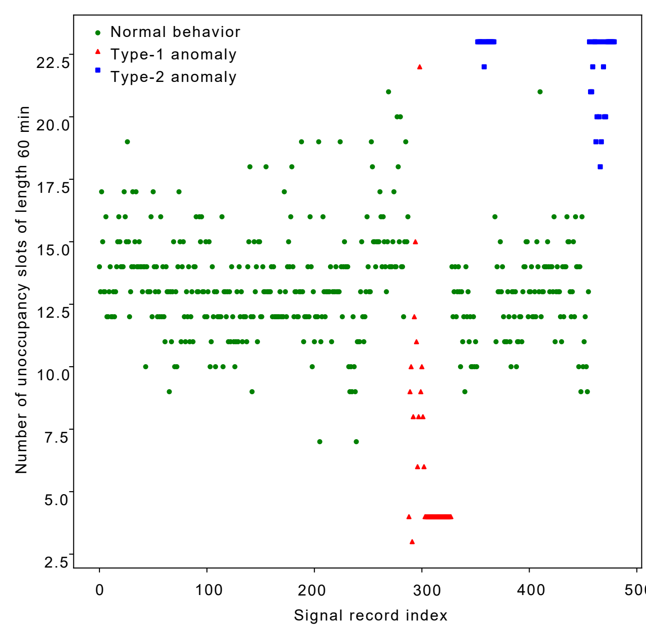

Fig. 3: Time slots of unoccupancy duration of length 60 min.

2) Frequency domain features: The discrete Fourier transform (DFT) of signal record $\{x_i\}$ is given by

$$X_{pi} = \sum_{n=0}^{N-1} x_i[n]. \exp^{-\frac{2\pi j}{N} pn}.$$

We consider the following features related to the energy in low, medium and high frequency bands:

$$F_i^{low} = \sum_{p=0}^{\lfloor N/6 \rfloor} |X_{pi}|^2, \quad (4)$$

$$F_i^{med} = \sum_{p=\lfloor N/6 \rfloor+1}^{\lfloor N/3 \rfloor} |X_{pi}|^2, \quad (5)$$

$$F_i^{high} = \sum_{p=\lfloor N/3 \rfloor+1}^{\lfloor N/2 \rfloor} |X_{pi}|^2. \quad (6)$$

We plot F_i^{med} for different signal records, each record taken over a day, in Fig. 4. Signal records with type-1 anomaly generally have high energy content in the medium frequency band, while signal records with type-2 anomaly have lower energy content.

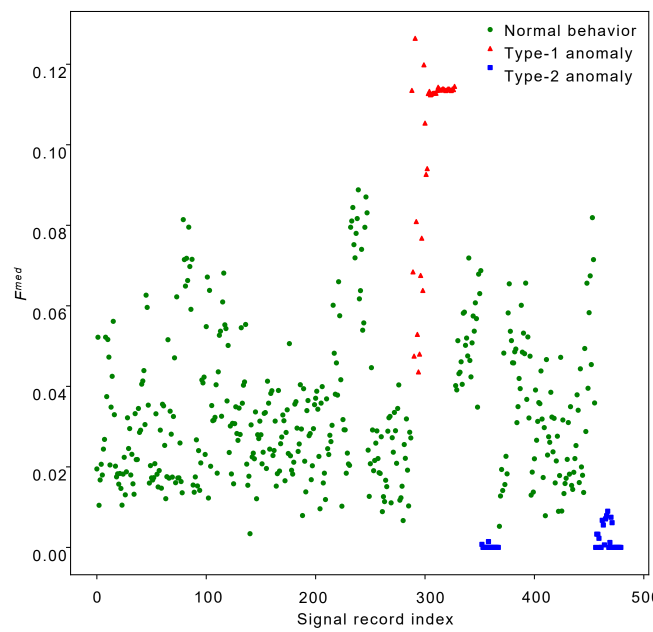

Fig. 4: F_i^{med}.

B. Random forest classifier

The proposed machine learning method to detect anomalous sensor behavior employs a 3-class random forest classifier trained with the set of features described in Section III-A, denoted as $\mathscr{V}_i = \{E_i^b, E_i^{b'}, u_i^{L_1}, u_i^{L_2}, F_i^{low}, F_i^{med}, F_i^{high}\}$, with $i \in$ training set S_{train}. Here L_1 and L_2 are two choices of length of unoccupancy slots (short and long respectively) in (3).

As almost all learning approaches, the proposed method consists into two phases: the *training* phase and the *test* phase. During the training phase a random forest applies the general technique of *bootstrap aggregating*, or *bagging*, to tree learners. In this way, each decision tree is fit using only a random subset of sensor signal records selected with replacement from the dataset. Random input selection not only improves classification accuracy but also increases the algorithm's speed [16]. The second source of randomness comes from the selection of a random subset of the features at each split inside the decision tree. In particular, each node of the tree implements a binary decision conditioned only on a randomly selected feature, and based on this decision the incoming dataset is split into two parts [17] [18].

An advantage of the random forest classifier lies in the limited number of parameters that need to be tuned. We need to choose the number of decision trees $\mathscr{T}$ composing the forest and their maximum depth. During the training phase, a cross validation technique is used to choose these two parameters. In particular, we run the K-fold cross validation: the training set is split into K subsets of equal size. Then, for each fold, the random forest is trained with different values for the parameters on the union of all the remaining subsets. The error of its classification output is computed on the fold which was left out. The chosen parameters will be the ones which yield the lowest error. Once the cross validation procedure is concluded, the algorithm is trained again using the tuned parameters on the whole training dataset.

In general we expect that signal records with anomalous behavior are much fewer in number than signal records with normal behavior. This means that the dataset on which the random forest is trained is imbalanced towards the normal behavior class. Since the performance of random forest classifiers is sensitive to imbalanced datasets [19], a possible way to deal with this problem is to set class weights such that they are inversely proportional to class frequencies on the bootstrap sample of the input data S_{train}, with which every tree is grown. Let us define the following notations:

- $\mathscr{C}_l$ is a particular class;
- w_l is the weight associated to class $\mathscr{C}_l$;
- $|\cdot|$ represents the cardinality of the argument;
- B_s, with $s \in \{1, \ldots, \mathscr{T}\}$ are the built bootstrapped samples from the original training set S_{train}, where $|B_s| = |S_{train}|$ and examples are chosen randomly with replacement from S_{train};
- $\mathscr{W}_{l,s}$ is the cardinality of class $\mathscr{C}_l$ inside B_s.

For class $\mathscr{C}_l$, the proposed method uses the following class weight:

$$w_{l,s} = \frac{|S_{train}|}{3 \cdot \mathscr{W}_{l,s}}, \qquad s = \{1, \ldots, \mathscr{T}\}, \tag{7}$$

where the factor 3 in the denominator is the number of classes.

The random forest algorithm also provides the relative importance of each feature as an output after the training phase [20]. Indeed, at each split inside the tree, the algorithm measures the feature importance by looking at how much the tree nodes, which use that feature, reduce impurity across all trees in the forest and it normalized the results, so that the sum of all importance is equal to 1. By looking at the feature importance, it is possible to decide which features may be dropped, or whether there is a need to engineer new features.

After the proposed algorithm has been trained, its performance is evaluated on the test set S_{test}: during the test phase, the set of features $\mathscr{V}_i$, $i \in S_{test}$ are computed for each sensor signal record belonging to S_{test} and used as input for the trained random forest classifier. The output of the model are the labels of the class to which each test record belongs to.

IV. Performance evaluation

We consider occupancy sensor signal data from a connected lighting system in an office building. Occupancy detection is generated every minute. Classification is performed at the end of each day (with $N = 24 \times 60 = 1440$). The training data consists of 750 signal records with normal sensor behavior, 25 signal records with type-1 anomaly and 75 signal records with type-2 anomaly. The test data consists of 400 signal records with normal sensor behavior, 40 signal records each for type-1 anomaly and type-2 anomaly.

We use an unsupervised k-means clustering method (with $k = 3$ classes) to showcase the performance improvements offered using a supervised random classifier approach. Furthermore, the improvements using multiple features are shown by comparing the proposed method with a random forest classifier with a single feature. The feature used in both k-means and the random forest classifier is the daily signal energy. At the end of each business day, a classification decision is made.

Based on inspection of the occupancy sensor signals, the business hour period is taken to be Monday to Friday between 7:00 am and 7:00 pm; the remainder is consider to be non-business hour period.

In the proposed method, features in (3) are computed considering unoccupied slots of duration 10 minutes and 60 minutes.

For comparison, we consider accuracy η, true positive rate (TPR) and the false positive rate (FPR) as performance measures. Denote the true class label of signal record $\{x_i\}$ by c_i and the estimated label by $\hat{c}_i$. For a particular class $\mathscr{C}_l$,

$$\eta = \frac{\sum_{l,i}[\hat{c}_i \in \mathscr{C}_l | c_i \in \mathscr{C}_l]}{\sum_{l,i} c_i \in \mathscr{C}_l} \tag{8}$$

$$\text{TPR}(\mathscr{C}_l) = \frac{\sum_{i}[\hat{c}_i \in \mathscr{C}_l | c_i \in \mathscr{C}_l]}{\sum_{i} c_i \in \mathscr{C}_l} \tag{9}$$

$$\text{FPR}(\mathscr{C}_l) = \frac{\sum_{i}[\hat{c}_i \in \mathscr{C}_l | c_i \notin \mathscr{C}_l]}{\sum_{i} c_i \notin \mathscr{C}_l} \tag{10}$$

The confusion matrix for the k-means method, random classifier with signal energy feature, and the proposed random forest classifier with multiple features is shown in Table I. The performance metrics η, TPR and FPR for the three methods are shown in Table II. The k-means method has relatively low accuracy, largely due to its poor ability to distinguish between sensors with normal behavior and ones with Type-1 anomaly. The random forest method with daily signal energy feature has reasonable performance, but fails to properly classify Type-1 anomalies. Specifically, this method incorrectly classifies many of the Type-1 anomalies as normal behavior as evidenced by a low TPR value of 10%. The proposed method clearly outperforms the two benchmark methods across all the classification

		Estimated label		
		Normal behavior	Type-1 anomaly	Type-2 anomaly
True label	Normal behavior	273 / 395 / 400	1 / 5 / 0	126 / 0 / 0
	Type-1 anomaly	31 / 36 / 6	4 / 4 / 34	5 / 0 / 0
	Type-2 anomaly	0 / 0 / 0	0 / 0 / 0	40 / 40 / 40

TABLE I: Confusion matrix (k-means/random forest with single feature/proposed random forest with multiple features).

	k-means	Random forest (one feature)	Random forest (proposed)
η	66.04	91.46	98.75
TPR(normal behavior)	68.25	98.75	100
FPR(normal behavior)	38.75	45	7.5
TPR(Type-1 anomaly)	10	10	85
FPR(Type-1 anomaly)	0.22	1.14	0
TPR(Type-2 anomaly)	100	100	100
FPR(Type-2 anomaly)	29.77	0	0

TABLE II: Performance of k-means, random forest with single feature, and proposed random forest with multiple features.

performance metrics. For type-2 anomalies, all three methods have a TPR of 100%, i.e. can easily classify Type-2 anomalies. This is since the dataset contained sensors with very sparse detection behavior.

V. Conclusions and discussions

We considered the problem of detecting anomalies in occupancy sensor detection signals in an office building environment where motion sensors are part of a connected lighting system. Temporal and frequency domain features were used in a random forest classifier to decide whether an occupancy sensor behavior was normal or exhibited one of two types of anomalies. We evaluated our proposed method using one month of occupancy sensor data from occupancy sensors in cell offices and meeting rooms. The detection accuracy was 98.75% on the test dataset, with a TPR of 85% and 100% for Type-1 and Type-2 anomaly respectively, and FPR of 0% in both cases.

In this paper, we made assumptions on business hours. In further work, we plan to derive this information from the occupancy sensor data itself. While the proposed method was a supervised machine learning method, a semi-supervised approach where sensors are assumed to be normally functional in the training phase will be investigated in next steps. We will also explore the use of spatio-temporal features across occupancy sensors as a possible way to improve the TPR of type-1 anomalies and FPR of normal behavior. We also plan to extend our evaluation to a larger dataset both in time and larger number of rooms, along with using data from more buildings.

References

[1] D. Caicedo and A. Pandharipande, "Sensor-driven lighting control with illumination and dimming constraints," *IEEE Sensors Journal*, vol. 9, no. 9, pp. 5169–5176, 15 2015.

[2] A. Pandharipande and D. Caicedo, "Smart indoor lighting systems with luminaire-based sensing: A review of lighting control approaches," *Energy and Buildings*, vol. 104, pp. 369 – 377, 2015.

[3] M.-S. Pan, L.-W. Yeh, Y.-A. Chen, Y.-H. Lin, and Y.-C. Tseng, "A WSN-based intelligent light control system considering user activities and profiles," *IEEE Sensors Journal*, pp. 1710–1721, 2008.

[4] G. R. Newsham, M. B. C. Aries, S. Mancini, and G. Faye, "Individual control of electric lighting in a daylit space," *Lighting Research & Technology*, vol. 40, no. 1, pp. 25–41, 2008.

[5] T. Hastie, R. Tibshirani, and J. Friedman, *The elements of statistical learning: data mining, inference and prediction*, Springer, Second edition, 2009.

[6] Y-J. Wen and AM Agogino, "Control of wireless-networked lighting in open-plan offices," *Lighting Research and Technology*, vol. 43, no. 2, pp. 235–248, 2011.

[7] A. Pandharipande, D. Caicedo, and X. Wang, "Sensor-driven wireless lighting control: System solutions and services for intelligent buildings," *IEEE Sensors Journal*, vol. 14, no. 12, pp. 4207–4215, 2014.

[8] D. Caicedo and A. Pandharipande, "Sensor data-driven lighting energy performance prediction," *IEEE Sensors Journal*, vol. 16, no. 16, pp. 6397–6405, Aug 2016.

[9] T. Labeodan, W. Zeiler, G. Boxem, and Y. Zhao, "Occupancy measurement in commercial office buildings for demand-driven control applications a survey and detection system evaluation," *Energy and Buildings*, vol. 93, pp. 303 – 314, 2015.

[10] C. Liao and P. Barooah, "An integrated approach to occupancy modeling and estimation in commercial buildings," in *Proceedings of the 2010 American Control Conference*, June 2010, pp. 3130–3135.

[11] V. L. Erickson, M. Á. Carreira-Perpiñán, and A. E. Cerpa, "Occupancy modeling and prediction for building energy management," *ACM Trans. Sen. Netw.*, vol. 10, no. 3, pp. 42:1–42:28, May 2014.

[12] A. Pandharipande, M. Zhao, E. Frimout, and P. Thijssen, "IoT lighting: Towards a connected building eco-system," in *IEEE 4th World Forum on Internet of Things (WF-IoT)*, Feb 2018, pp. 664–669.

[13] A. Pandharipande and G. R. Newsham, "Lighting controls: Evolution and revolution," *Lighting Research & Technology*, vol. 50, no. 1, pp. 115–128, 2018.

[14] B. Gunay and W. Shen, "Connected and distributed sensing in buildings: Improving operation and maintenance," *IEEE Systems, Man, and Cybernetics Magazine*, vol. 3, no. 4, pp. 27–34, Oct 2017.

[15] S. Li, A. Pandharipande, B. Masini, and D. Caicedo, "Automated detection of commissioning changes in connected lighting systems," *IEEE Internet of Things Journal*, pp. 1–1, 2018.

[16] L. Breiman, "Out-of-bag estimation," https://www.stat.berkeley.edu/~breiman/OOBestimation.pdf.

[17] T. G. Dietterich, "An experimental comparison of three methods for constructing ensembles of decision trees: Bagging, boosting, and randomization," *Machine Learning*, vol. 40, no. 2, pp. 139–157, Aug 2000.

[18] T. M. Khoshgoftaar, M. Golawala, and J. Van Hulse, "An empirical study of learning from imbalanced data using random forest," in *Proceedings of the 19th IEEE International Conference on Tools with Artificial Intelligence - Volume 02*, 2007, pp. 310–317.

[19] Y. Sun, *Cost-sensitive Boosting for Classification of Imbalanced Data*, Ph.D. thesis, Waterloo, Ont., Canada, Canada, 2007.

[20] G. Louppe, L. Wehenkel, A. Sutera, and P. Geurts, "Understanding variable importances in forests of randomized trees," in *Advances in Neural Information Processing Systems*.