Activity Recognition using Multi-Class Classification inside an Educational Building

Anooshmita Das*, and Mikkel Baun Kjærgaard†
Center for Energy Informatics, University of Southern Denmark, Odense, Denmark
Email:*adas@mmmi.sdu.dk ‡mbkj@mmmi.sdu.dk

***Abstract*—Activity Recognition can be referred to as the process of describing and classifying actions, pinpoint specific movements, and extract unique patterns from the dataset using heterogeneous sensing modalities. Activity Recognition approaches have garnered the attention of researchers in the energy management domain to enhance energy utilization in buildings. In our experiment, we define activities as a combination of different actions, which are detected using multiple sensors. To learn insights for the various activities, we used inexpensive Passive Infrared (PIR) sensors in the test-bed. This study aims at gaining high-level knowledge about activities from the low-resolution sensors deployed. For accurate occupancy counts, we have used 3D Stereo Vision Cameras at the entrance, and count lines are defined to capture the transitions of inflow and outflow of multiple occupants. Multi-class labels enable activity recognition on the collected dataset. The multi-class labels used are 1) Moving, 2) Stagnant, 3) Outside, 4) Both (Moving and Stagnant), 5) No activity inside. The labeling for the multi-class is done through an algorithm using supervised learning. The data acquisition gets carried out from 23rd November to 3rd December 2018, spanning over a period for 11 days. The results document that Gradient Boosting Classifier outperforms any other Machine Learning Classification (MLC) algorithm with an accuracy of 97.59% and an F1 score of 97.40% for activity recognition. This paper also explicitly highlights the challenges and limitations faced during the initial phase for the deployment, and it identifies the key research trends and directs towards the potential improvements in the field of occupancy sensing for energy-efficient buildings.**

***Index Terms*—Building Performance, Occupant Behaviour, Sensor Fusion, Pattern Recognition, Knowledge Discovery, Machine Learning, Activity Recognition**

I. Introduction

With the advancement of ubiquitous computing and the availability of a wide range of low-cost sensing modalities, human activity recognition (HAR) has garnered attention from both researchers and industry [1]–[5]. AR also has many notable and compelling applications ranging from healthcare to fitness, elderly assistance, ambient sensing, indoor navigation, gaming, are to name a few [2], [3], [6], [7]. Activity Recognition (AR) using various sensing modalities has emerged as a game-changer for automated monitoring of occupant behavior analysis [6], [8].

The primary goal of AR is modeling and making informed decisions about occupant behavior and actions along-with automated identification of the tasks in real-world settings by processing the spatial and temporal features acquired from diverse sensing modalities [6]. The massive data generated from the building systems contain valuable and rich information that requires to be mined and facilitate up-to-date actions and informed decision making. Generally, the execution of different daily activities by occupants gets performed under diverse contexts, for example- specific scenario, location, time, and object [9], [10].

In this paper, Activity Recognition (AR) can be defined as pin-pointing specific movements and learning profound knowledge about the actions of multiple occupants using low-level sensing modalities inside a monitored space [6]. AR is also concerned with the assignment of an activity label to a sequence of sensor events generated from the heterogeneous sensors available from building systems. However, in real-world applications, most of the data remain unlabelled (unsupervised), and activities take place continuously, i.e., occupants quickly transfer from one activity to another. Activity recognition should aid context-aware computing to handle Spatio-temporal data. Activity Recognition method implemented should be flexible to capture the variations and diversity of occupant behavior and actions.

In this research work, an educational building gets considered for the experiments, and the occupants complete their mundane tasks without any intervention in a no-restriction environment. Activity Recognition is a pivotal process to incorporate ambient intelligence into a smart environment. AR includes a complex process of active monitoring, sensor mapping, activity modeling, reasoning, making informed decisions, and inferences about an activity. It is imperative to recognize the ongoing activity and identify exposed abnormalities, extract contextual attributes related to that activity and automate actions within the monitored area. The mundane tasks are the frequent actions that are repeated by the occupants and these routine actions or activities are considered as it helps to set a defined boundary for the daily activities performed by the occupants. The most challenging element of multi-occupant activity recognition is data association, i.e., mapping an occupant with what caused the sensor triggering while using non-intrusive sensors that cannot directly identify the individual [6].

The stochastic and dynamic nature of occupancy behavior increases the complexity of identifying the activities precisely [6]. Therefore, sensor fusion seems a feasible option to leverage the unique capabilities and overcome the limitations for each sensing modality [11]. This paper presents the sensor fusion of inexpensive Passive Infrared (PIR) sensors and 3D

Stereo Vision Camera to detect the activities within the context of training, testing and validation dataset.

Cameras are considered intrusive; therefore, we have followed protocols and asked the Building Management System (BMS) authorities for proper authorization and permission. The recordings from 3D Stereo Vision Cameras were obtained for ground truth validation; the recordings were procured with stringent compliance to user's privacy and Danish Regulation to collect scientific data. The occupant confidentiality is preserved, and the video recordings were deleted after the validation was done. For further evaluation, a confusion matrix is created to infer different underlying causes for the discrepancies of imbalanced class distribution within the dataset. The data is collected for the entire test period of 11 days spanning from 23rd November to 3rd December 2018. The dataset used for this experiment is good enough to capture the semantic occupancy presence impression and extrapolate unique information about occupancy behavior and actions.

The significant goals considered in this paper gets described below-

1) Propose a methodology to transform the time-series data from heterogeneous sensor readings into feature vectors to have labeled supervised learning.

2) Dataset gets split into training, testing, and validation set to evaluate the feature vector using cross-validation techniques such as - 10 Fold Cross-Validation and evaluation metrics such as - accuracy, F1 Score, and misclassification rate.

3) Comparative analysis and inference of performance for different Machine Learning Classification (MLC) algorithms for identifying activities.

4) Comprehensive discussion for decoding the significant challenges of multi-occupant activity recognition and revealing insights on dynamic occupant behavior.

The significant contributions of this paper are:

1) Propose a label based approach for activity recognition using traditional and advanced Machine Learning Classification (MLC) algorithms. It highlights the experimental results, which MLC algorithm provides better results using k-fold cross-validation techniques. We evaluated it using metrics such as accuracy, F1 Score, and misclassification rate.

2) Crafted an experiment to demonstrate a real case scenario deployment and proposed a method for accurately estimating occupant activities using various sensing modalities in a non-intrusive and reliable manner.

The rest of the paper gets organized as follows: Section 2 gives an overview of the related work. Section 3 provides an overview of the case scenario. Section 4 describes the methodologies used for data alignment, data cleaning, pre-processing, sensor fusion, and data aggregation. It also elucidates the data labeling algorithm for labeling the instances for multi-class classification. Section 5 discusses the evaluation and results. Section 6 provides a discussion and highlights the limitations and future scope of the proposed activity recognition method. In Section 7, a conclusion gets drawn.

II. Related Work

The growth and improvement in wireless technologies and Internet of Things (IoT), augmented with Machine Learning (ML) and Deep Learning (DL) paradigms, the smart environment gets transformed and revolutionized. Diverse applications can be underpinned by using state-of-the-art machine learning algorithms such as - efficient space utilization, intelligent building operations, improved heating, ventilation and air conditioning (HVAC) conditions, safety, and evacuation, to name a few. [12], [13]. Recently, researchers have witnessed an increasing focus in the prominence of activity recognition for multiple occupants in offices, study-rooms, meeting-rooms, and public places to study the occupancy behavior influence on energy consumption and management domain inside a building [6]. Activity recognition has also garnered the attention of researchers as it has other notable applications such as healthcare, lifestyle monitoring, elderly assistance, and indoor localization. Traditional Machine ML methods such as Support Vector Machines (SVM), Decision Trees (DT), Naive Bayes, Markov Models have significantly fostered the research of activity recognition during the past few decades. The recent works had been summarized in Table I.

III. Experimental Set-Up

In this paper, an experiment is crafted in a large study-room in an educational institute. We deployed 30 PIR sensors and 2 3D Stereo Vision Cameras from Xovis manufacturers, at the entrance/exit points to capture the transitions of occupants inside-out through count lines. The test space is occupied by multi-occupants performing their mundane tasks in a no-restriction environment. The PIR sensors used for this case was Hamilton H7C. The PIR sensors are affordable, easy to deploy with a longer battery-life. The detection angle for Field of View(FoV) of Hamilton was 43-47 °and detection range is 5 meters. The sampling frequency of the PIR sensors is 20 seconds. In this case, the sensors are mounted at a height of 3.8 meters. These Hamilton's are designed to have a battery-life of up-to 5 years. Figure 1 shows the sensor placement layout for the deployed case scenario. The sensor can capture motion and presence, air humidity, illuminance, acceleration and temperature. The Hamilton sensor connects to the cloud via a gateway and data is extracted via a cloud-based REST API provided by the Hamiltoniot. Each PIR has an unique *sensor-id*, represented by green dots in Figure 1.

IV. Methodology

The overview of the proposed activity recognition method is given in Figure 2 . The crucial components are listed below:

A. Time Series Data: Acquisition

The time-series data (in A) gets collected from heterogeneous sensing modalities such as - PIR sensors and 3D Stereo Vision Camera. Once the data is gathered, pre-processing (in B) is a crucial step implemented to have a structured dataset from the raw sensor readings, see Figure 2.

TABLE I
OVERVIEW OF RELATED WORK

Author and Year	Sensors Used	Benchmark Dataset	Building Type	Targeted Activities	Characteristics	Model Used	Evaluation
Koping et. al 2018 [14]	Inertial measurement units (Smart-phones smart-watches, and smart-glasses)	Empirical Data	Office	11 different activities, out of which 4 were static (lying, sitting, standing, and walking) and 7 were dynamic (bending, getting up, lying down, putting a hand back, sitting down, standing up, and stretching a hand)	The wearable smart devices capture the data from different sensors, and each type of sensor data is used to construct a codebook. An SVM is constructed for each of the 11 activities.	SVM	Accuracy of 87.1% with the SVM model, where the evaluations for static and dynamic activities are separated Mean of Average Precisions (MAP) for the 11 activities was 88%.
Bulling et. al 2014 [15]	Inertial measurement units	Empirical Data	Office	11 activities: Gestures for common use and playing tennis	IMUs record timestamped motion data at a sampling rate of 32Hz. Data Synchronization was performed using SenseHub	DA, NB SVM, HMM JB, and k-NN	Precision and Recall (person-dependent) DA: 87.3%, 70.6% NB: 78.2%, 69.2% SVM: 96%, 84.8% HMM: 83.6%, 85.6% JB: 81.3%, 82.6% k-nn: 94.1%, 62.4%
Edel et. al 2016 [16]	Inertial measurement units	1.Opportunity 2.PAMAP2 3.Custom	NA	Opportunity dataset contains human activities recorded from on-body sensors from 4 participants. The PAMAP2 dataset contains data of 18 different physical activities, performed by 9 participants located on the hand, chest and ankle over and a temperature and heart rate sensor. Custom Dataset has 22 participants, performed 15 times each, one of the 14 different activities.	Demonstrated the suitability of using a binary representation of the BLSTM recurrent neural network	Binarized BLSTM-RNN	The binarized BLSTM-RNN achieves almost every time the second best performance in every configuration
Rawashdeh et. al 2017 [17]	Microsoft Kinect	1.Cairo 2.Milan 3.Tulum	Smart Home	Cairo dataset- 10 activities Milan dataset-15 activities Tulum Dataset- 16 activities	The induction of the fingerprint features, as an extra dimension to recognize activities , played a significant role in minimizing the false positive rate; i.e. minimize mis-classified instances	Naive Bayes, SVM, and J48 Classifiers	ANOVA statistical testing to measure whether the difference in terms of TP, FP, and F-measure is significant or not. Results show that at(p<0.01) confidence-interval, the differences were statistically significant.

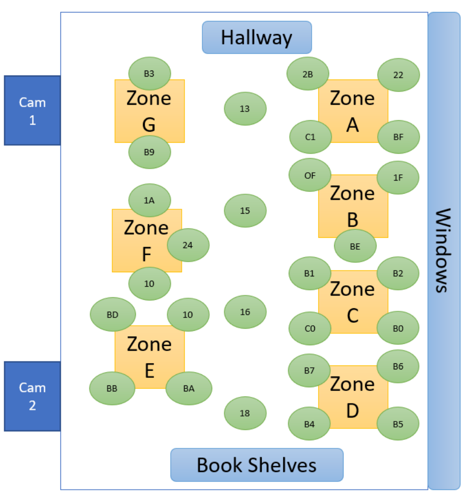

Fig. 1. Case Scenario OU44 Study Zone: Sensor placement PIR (green dots) and Cameras (Entrance/Exit). The yellow elements are the tables and the light blue elements are the furniture

a) Dataset: We used a dataset from a large office building, 8000 m^2, mainly comprising of office-rooms, classrooms, and study zones. The room type considered for this experiment was a **study zone**. *The dataset contains 5011 PIR readings, 3847 people counts*. The people counting camera from Xovis is placed at a strategic point to cover the transitions flow of occupants through the entrances and exits of the study zone through count lines. We used a dataset spanning from from 23rd November to 3rd December 2018. Refer to Table 2; we can see an example for the dataset.

Later, we want to release the full dataset which can serve for benchmarking and foster data-driven research in occupant activity recognition. This dataset used in this paper would help to extract the inter-relationships of different factors influencing occupant behavior and actions. *The authors in this paper also released an open dataset based on previous work for extracting occupancy presence patterns and trajectories* [18].

TABLE II
DATASET EXAMPLE

Timestamps	PIR 1	PIR 2	...	PIR30	Forward_Count (Xovis Camera)	Backward_Count (Xovis Camera)	Labels
t0	0	0	...	0	0	0	No_Activity
t1	0	1	...	1	1	0	Moving
...	...	...	...	...	...	...	...
tn	0	0	...	0	0	1	Outside

B. Pre-processing

Firstly, the training data is aligned based on the timestamps using the pandas library. As part of the pre-processing, data cleaning is performed to remove inconsistencies and missing data in the dataset. This step ensures that the data

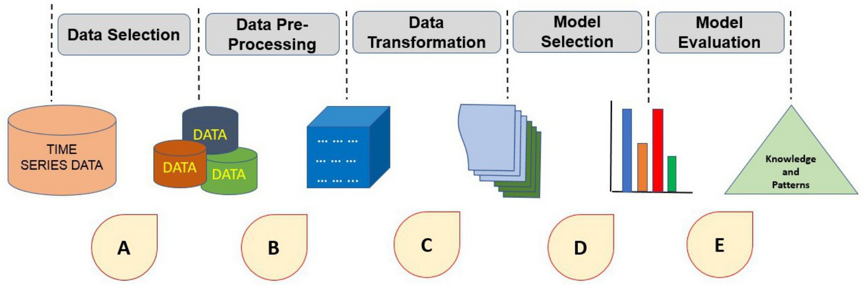

Fig. 2. Overview

is relevant, and it can be transformed to fit into the model framework. Up-sampling is performed at a 1-minute granularity because different sensors had varying sampling rates. PIR sensor triggers every 20 seconds and the camera resolution was 1 minute. So we combined 3 PIR samples to match the granularity of the camera. Data normalization (min-max normalization) was performed to reduce the redundancy, and it guarantees a uniform and structured dataset to work for further analysis. It helps to avoid the dependence on the specific choice of measured units instead provide equal weight to all the attributes in their representation in the dataset.

1) Training, classification, and Testing based on feature vectors: The fundamental task of this research work recognized the activities of multi-occupants based on the test dataset. The entire dataset gets split into 7 days of training data consisting of labeled instances and 4 days of the unlabeled test dataset. A *feature* can refer to as a function of heterogeneous sensor data over a fixed time duration, and a *class* can be defined as the representation of occupancy state and action at any given instant of time.

Algorithm 1 Labeling the Instances: Supervised Learning

Input : Time-Series Data combined in a CSV file which includes Xovis-Cameras and PIR readings
Output: Labelled instances
A list is created consisting of PIR and Xovis Camera readings
list = [$\mathbf{PIR_1}$... $\mathbf{PIR_{30}}$, Xovis-Camera 1, Xovis-Camera 2]
if *No Camera and PIR sensor activates* **then**
| **No-Activity Inside label** is assigned to that instance
else if *Only xovis camera activates* **then**
| **Outside label** is assigned to that instance
else if *trigger-count of any PIR $\geq$ 2 and other PIR sensors = 0* **then**
| **Stagnant label** is assigned to that instance
else if *trigger-count of any PIR $\geq$ 2* ***and*** *at-least one PIR trigger-count in different zone = 1* **then**
| **Both label** is assigned to that instance
else
| **Moving label** is assigned to that instance
end

Before splitting the data, we performed data annotation on the dataset. It is common practice to attach a label (annotation) that is representing the activity in which a given event belongs to, and is known as Data Annotation. **Algorithm 1** shows the logic behind labeling the instances in the training dataset. The multi-class classification labels used are 1) **Moving**- *a person is showing specific movements or walking within the test space*, 2) **Stagnant** - *a person is still inside the test space, for example - sitting*, 3) **Outside** - *a person who is outside the vicinity of the test space, but crossed the count line of the camera, without entering the study zone*, 4) **Both (Moving and Stagnant)** - *a single sensor is constantly triggering which means a person is still at that position, also other sensors are triggering at the same time which indicates movement in a different zone', we can assume that multiple occupants are present in the test bed*, 5) **No activity inside** - *if there is no sensor triggering for a prolonged duration inside the test space. We consider that there is no occupant performing activities*, this is important to consider because after the office hours or during the night time, there are no occupant's inside the test space. The testing and evaluation get explicitly described in Section 5.

2) Privacy and Data Suppression:: The most sensitive part of this dataset is the data collected outside the opening hours inside the public space since those readings get plausibly collected from employees working within the monitored area. This part of the data can be used by the employer to administrate and estimate the work performance of the employees in the area. Due to this, we do not release any of the sample readings collected outside of the opening hours [19]. Furthermore, the occupants in the monitored area also have the right to privacy. For the protection of the identity of the days, stringent measures get undertaken. We do not include the dates as part of the dataset. Also, re-ordering the days by mapping the date component between 0-10, thus creating a random permutation of the days. These precautions make it significantly more difficult for adversaries to perform data linkage attacks upon the released data, and hence identifying/revealing the identity or critical information of the occupants or the location of the monitored area. In the dataset, we have introduced a workday indicator that accounts for weekends and national holidays.

C. Sensor Fusion and Data Transformation

Sensor fusion intelligently combines data from the heterogeneous sensor to achieve enhanced accuracy and derive more contextual knowledge and extract unique patterns about occupant activities. Data aggregation merges the PIR readings and camera counts into a coherent structure (C) in Figure 2 and eliminates the noise and variability in the dataset.

D. Model Selection, Comparison and Evaluation

To precisely evaluate (in E) the performance and model comparison (see Figure 2), metrics such as accuracy, F1 score, misclassification rate get used for different MLC algorithm such as - Decision Trees (Adaboost) (DT), Random Forest (Adaboost) (RF), Gradient Boosting, K- Nearest Neighbor (KNN), Naive Bayes, Support Vector Machine (SVM), XG-Boost (in D) using a dataset from a large office building. For cross-validation, 10 fold cross-validation technique gets implemented.

V. Evaluation and Results

This section presents and discusses the results achieved after applying pre-processing and sampling techniques, feature selection methods, and implementing MLC algorithms to the dataset.

TABLE III
Activity Recognition Scores for different classifiers

Classifier	Accuracy	F1 Score	Misclassification
Decision Trees	91.08 %	89.75 %	0.085 %
Random Forests	91.25 %	90.02 %	0.063 %
Gradient Boosting	97.59 %	97.40 %	0.024 %
KNN	93.76 %	93.21 %	0.062 %
Naive Bayes	88.57 %	88.07 %	0.114 %
SVM	92.70 %	91.53 %	0.073 %
XGBoost	96.93 %	96.63 %	0.030 %

There are 231 labelled instances for **Moving**, 48 labelled instances for **Stagnant/Sitting**, 106 labelled instances for **Outside**, 287 labelled instances for **Both (Moving and Stagnant)** and 2281 labelled instances for **No activity inside** in the dataset, see Figure 3. As we deployed non-intrusive sensors, the experiment is not inclusive of all the complex activities. One of the major limitations of using non-intrusive sensors is data mapping with the underlying cause of what activated the sensor triggering.

Table 3 highlights the results for the different classifiers after implementing the evaluation metrics. It provides a comparison of accuracy, F1 Score, and Misclassification rate of different MLC algorithms. Gradient Boosting Classifier achieves an accuracy of 97.59 %, F1 Score of 97.40 % and misclassification rate of 0.024 %. Figure 4 shows the accuracy score for DT Adaboost, RF Adaboost, KNN, Naive Bayes, SVM, XGBoost, which is 91.08 %, 91.25 %, 93.76 %, 88.57 %, 92.70 %, 96.93 %, respectively. Figure 5 shows the F1 score for DT Adaboost, RF Adaboost, KNN, Naive Bayes, SVM, XGBoost, which is 89.75 %, 90.02 %, 93.21 %, 88.07 %, 91.53 %, 96.63 %, respectively.

For the ground truth, count lines are defined to monitor the transitions (entrances/exits) of occupancy flow. We also have procured video recordings for the test-period, however validation of the ground truth data was tedious and manually expensive.

VI. Discussion, Limitation and Future Scope

The primary purpose of this research is the design and development of an affordable and non-intrusive solution to detect occupant activities inside a monitored space using inexpensive PIR sensors. Choosing the right sensor is critical for successful activity recognition. The sensor placement strategy is another critical criteria to recognize fine-grained activities accurately.

We also identify that there is a requirement to classify or differentiate the sub-activities instead of assignment of the same labels to similar activities or likely to be similar activities. Besides, the pre-processing might have added irrelevant data, which may have affected the performance of the MLC

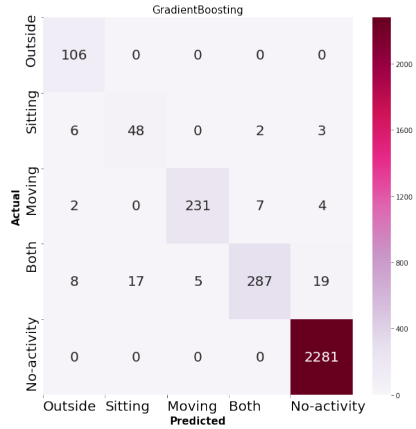

Fig. 3. Confusion Matrix

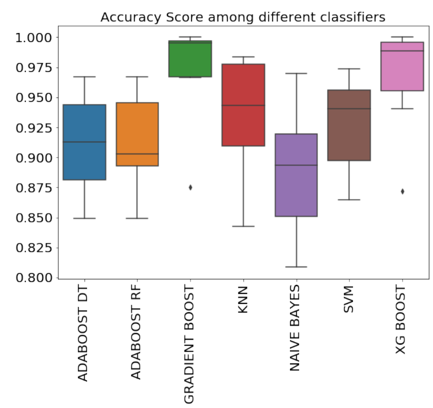

Fig. 4. Accuracy Score for Different Classifiers

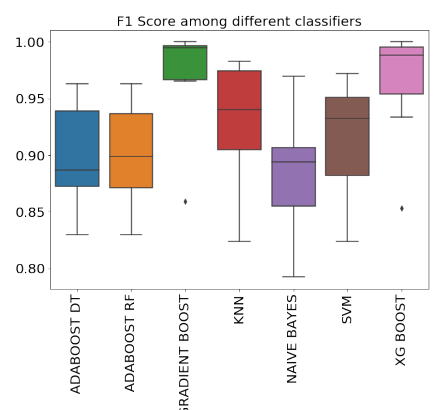

Fig. 5. F1 Score for Different Classifiers

algorithms negatively. There is a limitation in this study that is not inclusive of all kinds of activities. Thus, our future work would focus on further investigation of deploying acoustic or wearable sensors to detect other kinds of activities.

Another plan of work would also require the creation of a knowledge base for the activities. That would help in the meta-data construction and mapping of the sensor to the activities. We would also like to explore and extend the study-period.

Significant advancement by deep learning methods is achieved in the field of activity recognition. It has also alleviated the curse of the dimensionality that traditional ML-based activity recognition methods suffer from. One of the limitations of ML algorithms applied in activity recognition is that the feature extraction relies upon manually designed features, domain knowledge, and experience, which hinders the generalization of the model. ML algorithms cannot handle complex activity recognition scenarios, which are tedious and manually expensive in terms of data collection, labeled annotations, and model construction. The first advantage of using Deep learning models is that feature extraction is efficient with little efforts from humans. Secondly, the deep learning models can get reused for similar problems, and the construction of the DL models is efficient.

Therefore, advancement in deep learning makes it possible for automated feature extraction and selection. The features can be learned automatically by the network instead of manual assignment. DL can be applied when there is a lack of domain knowledge for feature introspection. Shallow features can be recognized well with ML but a difficulty in recognizing context-aware activities (for example: working on laptop). In traditional approaches, extensive training data and labeled annotations are mandatory for supervised learning, but in real-world applications, most of the data remain unlabelled (unsupervised). Due to this, typical models are unadaptable to a diverse range of context-aware activity recognition configurations. In future work, DL can be conscripted in the application when the data size is large; for smaller datasets, ML is preferable. DL outshines/outperforms for complex prediction tasks. Unlike ML approaches, DL classifiers train through feature learning rather than task-specific algorithms. As part of future work, we want to explore deep learning models and compare the performance of activity recognition.

VII. Conclusion

Activity recognition is a field concerned with identifying specific movements of occupants based on heterogeneous sensor data. Annotated Labels get used for typical activities performed within the monitored space such as - walking, sitting, etc. Activity recognition based on sensor data requires a profound high level of knowledge about occupant activities using multitudes of low resolution sensing modalities. In this paper, we implemented and demonstrated the activity recognition using Machine Learning Classification (MLC) algorithms based on a real case scenario. The results also provide a performance comparison among different traditional and advanced MLC algorithms. The patterns extracted from the dataset can grant access to the new dimension of investigation associated with the dynamic occupant behavior inside the commercial/public/educational buildings.

References

[1] S. Zolfaghari and M. R. Keyvanpour, "Sarf: Smart activity recognition framework in ambient assisted living," in *FedCSIS'16*. IEEE, 2016, pp. 1435–1443.

[2] A. Murad and J.-Y. Pyun, "Deep recurrent neural networks for human activity recognition," *Sensors*, vol. 17, no. 11, p. 2556, 2017.

[3] D. Marikyan, S. Papagiannidis, and E. Alamanos, "A systematic review of the smart home literature: A user perspective," *Technological Forecasting and Social Change*, vol. 138, pp. 139–154, 2019.

[4] C. Wilson, T. Hargreaves, and R. Hauxwell-Baldwin, "Benefits and risks of smart home technologies," *Energy Policy*, vol. 103, pp. 72–83, 2017.

[5] M. Li, W. Gu, W. Chen, Y. He, Y. Wu, and Y. Zhang, "Smart home: Architecture, technologies and systems," *Procedia computer science*, vol. 131, pp. 393–400, 2018.

[6] A. Das, F. C. Sangogboye, E. S. K. Raun, and M. B. Kjærgaard, "Heterosense: An occupancy sensing framework for multi-class classification for activity recognition and trajectory detection," in *Proceedings of the Fourth International Workshop on Social Sensing*. ACM, 2019, pp. 12–17.

[7] J. Wen and Z. Wang, "Learning general model for activity recognition with limited labelled data," *Expert Systems with Applications*, vol. 74, pp. 19–28, 2017.

[8] M. U. Mehmood, D. Chun, Z. Zeeshan, H. Han, G. Jeon, and K. Chen, "A review of the applications of artificial intelligence and big data to buildings for energy-efficiency and a comfortable indoor living environment," *Energy and Buildings*, p. 109383, 2019.

[9] D. Riboni and C. Bettini, "Cosar: hybrid reasoning for context-aware activity recognition," *Personal and Ubiquitous Computing*, vol. 15, no. 3, pp. 271–289, 2011.

[10] D. Riboni, T. Sztyler, G. Civitarese, and H. Stuckenschmidt, "Unsupervised recognition of interleaved activities of daily living through ontological and probabilistic reasoning," in *Proceedings of the 2016 ACM International Joint Conference on Pervasive and Ubiquitous Computing*. ACM, 2016, pp. 1–12.

[11] A. Das, F. C. Sangogboye, and M. B. Kjærgaard, "Mapping the important sensor limitations to design robust occupant sensing systems," in *Proceedings of the 2018 ACM International Joint Conference and 2018 International Symposium on Pervasive and Ubiquitous Computing and Wearable Computers*. ACM, 2018, pp. 1156–1161.

[12] S. Ahmadi-Karvigh, A. Ghahramani, B. Becerik-Gerber, and L. Soibelman, "Real-time activity recognition for energy efficiency in buildings," *Applied energy*, vol. 211, pp. 146–160, 2018.

[13] A. Das and M. B. Kjærgaard, "Precept: occupancy presence prediction inside a commercial building," in *Adjunct Proceedings of the 2019 ACM International Joint Conference on Pervasive and Ubiquitous Computing and Proceedings of the 2019 ACM International Symposium on Wearable Computers*, 2019, pp. 486–491.

[14] L. Köping, K. Shirahama, and M. Grzegorzek, "A general framework for sensor-based human activity recognition," *Computers in biology and medicine*, vol. 95, pp. 248–260, 2018.

[15] A. Bulling, U. Blanke, and B. Schiele, "A tutorial on human activity recognition using body-worn inertial sensors," *ACM Computing Surveys (CSUR)*, vol. 46, no. 3, p. 33, 2014.

[16] M. Edel and E. Köppe, "Binarized-blstm-rnn based human activity recognition," in *2016 International Conference on Indoor Positioning and Indoor Navigation (IPIN)*. IEEE, 2016, pp. 1–7.

[17] M. Rawashdeh, M. G. Al Zamil, S. Samarah, M. S. Hossain, and G. Muhammad, "A knowledge-driven approach for activity recognition in smart homes based on activity profiling," *Future Generation Computer Systems*, 2017.

[18] A. Das, J. H. Schwee, E. S. Kolvig-Raun, and M. B. Kjærgaard, "Dataset: Occupancy presence and trajectory data from an instrumented public building," in *Proceedings of the 2nd Workshop on Data Acquisition To Analysis*, 2019, pp. 43–46.

[19] A. Das, E. S. K. Raun, and M. B. Kjærgaard, "Cam-stitch: Trajectory cavity stitching method for stereo vision cameras in a public building," in *Proceedings of the First International Workshop on Challenges in Artificial Intelligence and Machine Learning for Internet of Things*, 2019, pp. 8–14.