Communication-aided Terahertz Sensing: A Novel Indoor People Counting System Via Deep Learning

Eslam Hasan*1, Elmahedi Mahalal†2, Muhammad Ismail*3, Zi-Yang Wu‡4,
Mostafa M. Fouda§¶5, and Zubair Md Fadlullah‖6
*Department of Computer Science, Tennessee Technological University, Cookeville, TN, USA,
†Department of Computer Science, University of New Haven, West Haven, CT, USA,
‡College of Information Science and Engineering, Northeastern University, Shenyang, China,
§Department of Electrical and Computer Engineering, Idaho State University, Pocatello, ID, USA,
¶Center for Advanced Energy Studies (CAES), Idaho Falls, ID, USA,
‖Department of Computer Science, Western University, London, ON, Canada,
Emails: {1ebhasan42, 3mismail}@tntech.edu,2emahalal@newhaven.edu, 4wuziyang@ise.neu.edu.cn,
5mfouda@ieee.org, 6zfadlullah@ieee.org

***Abstract*—Indoor people counting systems are used in security monitoring, energy management, room resources adjustment, market research, and smart homes. However, the existing radio-frequency-based indoor people counting systems use a different frequency than the utilized radio frequency communication signal, which adds more costs for system deployment and wastes the radio frequency resources. This paper introduces a novel communication-aided terahertz (THz) sensing approach for robust indoor people counting that utilizes the THz communication downlink signal to sense the number of indoor people. Leveraging the wireless THz communication downlink signal, we propose a device-free, cost-effective, and non-intrusive indoor people counting system. Our method employs a 1D convolutional neural network (CNN) to process historical THz downlink channel gain data and accurately estimate the number of indoor occupants. The numerical results demonstrate the effectiveness of the proposed approach, achieving a remarkable 99.5% accuracy in people counting indoors up to eight people. The proposed model's ability to maintain high accuracy in indoor people counting across different numbers of users demonstrates its effectiveness and robustness in capturing the occupancy signature from the wireless THz downlink communication signal in indoor environments. Also, the accuracy of the proposed CNN time series classifier outperforms the random forest times series classifier with the catch22 feature extractor by more than 10% without needing any feature extraction methods. To the best of the authors' knowledge, this study represents the first investigation into indoor people counting in the THz frequency band utilizing the THz downlink communication signal for sensing the number of indoor occupants.**

***Index Terms*—Indoor people counting, terahertz, deep learning, communication-aided sensing.**

I. INTRODUCTION

Indoor people counting systems are increasingly used in applications such as security monitoring, energy management, room resources adjustment, market research, and smart homes [1]. Businesses benefit from this data by making informed decisions, such as optimizing resource allocation, predicting traffic trends, managing peak hours, and ensuring proper customer-staff ratios [2]. Additionally, it helps maintain occupancy capacity limits in real-time, promoting compliance with safety regulations like social distancing [3]. Also, indoor people counting systems are essential for smart buildings, enhancing their functionality and efficiency by providing real-time data on the exact number of people [4].

Indoor people counting systems enables resource allocation [5], safety [6], user experience enhancement [6], and energy efficiency [4]. For example, indoor people counting systems help optimize smart buildings' heating, ventilation, and air conditioning (HVAC) based on real-time occupancy information, leading to energy savings and efficiency [4].

Indoor people counting systems can be categorized into image-based and non-image-based systems [1]. Image-based systems typically use cameras and computer vision algorithms to count the number of people [7]. However, these systems suffer from privacy concerns [8]. On the other hand, non-image-based systems use technologies such as sensor-based or radio-frequency-based to count people [9]. Sensor-based methods involve various sensor combinations and require significant installation and maintenance expenses [4]. Hence, this paper focuses on non-image-based systems, especially radio-frequency-based, as they are cost-effective and non-intrusive [4].

The proposed method in [10] employs a support vector machine algorithm based on the wireless fidelity (WiFi) channel state information to count people. Ensemble learning is proposed in [11] using the WiFi signals for people counting. The proposed method in [12] utilizes a deep neural network based on impulse radio ultrawideband radar signals for people counting. Authors in [13] utilized a three-dimensional convolutional neural network (CNN) based on millimeter-wave (mmWave) signals for people counting. The proposed

This material is based upon work supported by the National Science Foundation (NSF) under Awards No. 2210251 and 2210252. Dataset generation was supported by the TNTech HPC cluster funded by NSF award No. 2127188.

method for people counting in [14] uses mmWave signals and combines multiple target tracking with a classifier to distinguish and count individuals within groups. Previous studies investigated people counting in frequency bands utilized in 5G networks and earlier. However, there is no investigation of people counting in higher frequencies, such as the Terahertz (THz) frequency band that will be utilized in 6G networks [15]. Also, the existing radio-frequency-based indoor people counting systems use a different frequency than the utilized radio frequency communication signal, which adds more costs for system deployment and wastes the radio frequency resources.

To fill this gap, we investigate people counting feasibility in the THz band. We propose an indoor people counting system that does not rely on cameras or additional sensors, which addresses privacy concerns and reduces installation and maintenance costs. We leverage the THz wireless network infrastructure for device-free, cost-effective, and non-intrusive indoor people counting. Hence, the THz wireless network will be used for both sensing and communication. To the best of our knowledge, this work represents the first investigation of the THz frequency band feasibility for counting people indoors. We validate the effectiveness of the proposed indoor people counting system through numerical simulations. The simulation results demonstrate the robustness and accuracy of the proposed indoor people counting system. The proposed indoor people counting system achieves a remarkable accuracy of 99.5% in counting up to eight people in the investigated indoor layout.

The rest of this paper is outlined as follows. Section II introduces the system model. Section III presents the proposed indoor people counting system. Section IV shows the numerical results. Section V concludes our findings.

II. System Model

This section outlines the indoor layout, the user mobility model, and the dynamic THz channel data generation.

A. Indoor Setup

We consider a standard meeting room with dimensions of 5m×5m×3m. The room is furnished with one meeting table and eight chairs arranged as shown in Fig. 1. Four THz access points (APs) installed on the ceiling support wireless communication within the room, as shown in Fig. 1. Mobile users within this room are modeled as cuboids with dimensions of 1.8m×0.2m×0.45m. These users have a weight of 70 kg and a speed of 2.1 m/s.

B. Mobility Model

This study uses a validated mobility model presented in [16] and used in the literature [17] to accurately simulate human movement within the indoor environment. The model operates on two timescales, which are macro and micro patterns. The macro pattern controls the user's movement decisions, including the timing and destination of their next move. It uses a semi-Markov renewal process incorporating regular return patterns and truncated Levy-walk behavior for unpredictable human motions. On the other hand, the micro pattern addresses the specifics of user movement, such as the shortest path to the destination, obstacle avoidance, and user equipment (UE) orientation. The mobility model is illustrated in Fig. 2. After generating the mobility traces, the user's position and the position of the serving THz AP are used to determine channel blockage. The THz channel gain is calculated based on the multi-ray THz channel propagation model described in the following subsection.

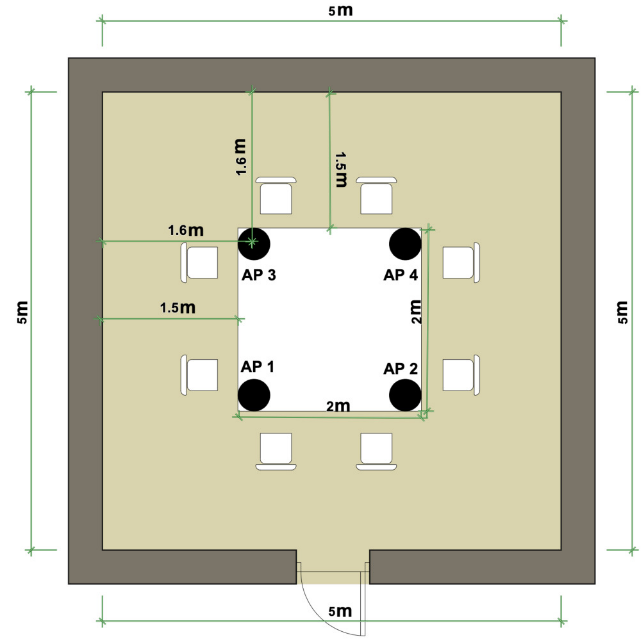

Fig. 1. Illustration of the indoor layout.

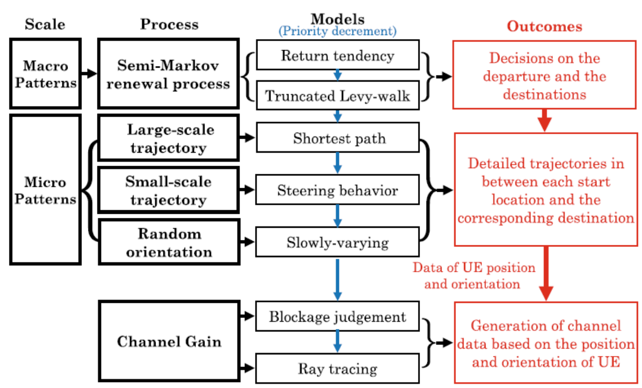

Fig. 2. Framework to generate dynamic THz channel gain.

C. Multi-ray Channel Propagation Model in the THz Band

This subsection outlines the multi-ray channel model that simulates the THz electromagnetic wave propagation using ray tracing techniques. The multi-ray channel model includes line-of-sight, reflected, scattered, and diffracted paths. However, diffracted paths can be disregarded in indoor THz environments since this phenomenon is significant only at lower microwave frequencies [18]. Additionally, scattered paths can be ignored due to the substantial losses they incur after scattering [18]. Therefore, only the line-of-sight (LoS) and reflection paths are considered. Hence, define the THz channel gain, $g(\tau)$, from the transmitter (AP) to the receiver (UE) as [19]

$$
\begin{aligned}
g(\tau) = &\left| \frac{c}{4\pi \cdot f \cdot r} e^{-\frac{1}{2}k(f)r} \right| \cdot \delta(\tau - \tau_{\text{LoS}}) \cdot 1_{\text{LoS}} \\
&+ \sum_{p=1}^{N_{\text{Ref}}} \left| \left(\frac{c}{4\pi f \cdot (r_1 + r_2)}\right) e^{-\frac{1}{2}k(f)(r_1+r_2)} \cdot \left(-e^{\frac{-2\cos(\theta_{\text{i}})}{\sqrt{n^2-1}}}\right) \right. \\
&\left. e^{-\frac{8\pi^2 \cdot f^2 \cdot \sigma^2 \cdot \cos^2(\theta_{\text{i}})}{c^2}} \right|_p \cdot \delta(\tau - \tau_{\text{Ref}}^p),
\end{aligned} \tag{1}
$$

where c is the speed of light, f is the frequency, r is the distance between the transmitter and the receiver, k is the medium absorption coefficient that depends on the utilized frequency, τ is the propagation delay, τ_{LoS} represents the LoS delay, 1_{LoS} is one if there is a line-of-sight path or zero otherwise, N_{Ref} is the number of reflected paths, r_1 is the distance between the transmitter and the reflector, r_2 is the distance between the reflector and the receiver, θ_{i} is the incident wave angle, n is the refractive index, σ is the rough surface height standard deviation, and τ_{Ref}^p is the delay of the p^{th} reflected path.

III. The Proposed Communication-aided THz Indoor People Counting

This section introduces a novel approach for indoor people counting, leveraging the unique properties of THz signals. First, we investigate the feasibility of indoor people counting in the THz band by examining the occupancy signature in the THz channel gain. Hence, we introduce the proposed indoor people counting system by formulating the problem of indoor people counting in the THz band, introducing the data preprocessing technique, and the proposed 1D CNN model training.

A. Occupancy Signature in THz Channel Gain

This subsection investigates the feasibility of indoor people counting in the THz band by examining the occupancy signature in the THz channel gain. Using the framework presented in Fig. 2, we simulated the downlink THz channel gain over 1000 mobility traces for one person, two, three, and eight people given the layout in Fig. 1.

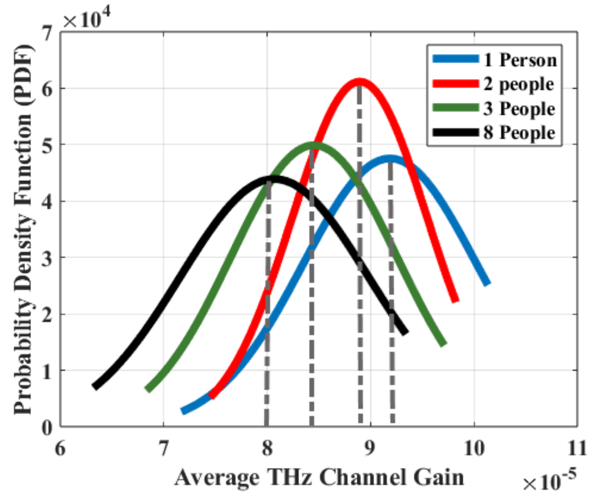

Fig. 3. The PDF vs. average THz channel gain.

Fig. 3 illustrates the probability density function (PDF) of the average THz channel gain for scenarios with one person, two, three, and eight people in the investigated indoor layout. In Fig. 3, the x-axis represents the average THz channel gain, while the y-axis represents the corresponding PDF. Fig. 3 demonstrates that each specific number of people follows a distinct distribution. Also, the pattern of the average THz channel gain in the investigated indoor room layout decreases with the increase in the number of people. This signature/pattern can be leveraged to develop algorithms for real-time monitoring of people count, where the THz channel gain data is continuously analyzed to determine the exact number of people.

B. Proposed People Counting System

Fig. 4 demonstrates the proposed indoor people counting system. The system input consists of historical channel gain data of the THz AP downlink signal. The input is preprocessed to ensure that the data fed into the CNN is of high quality and in a suitable format for analysis. The preprocessed data is then input into a 1D CNN. The 1D CNN is responsible for extracting occupancy information by identifying patterns and features in the input data correlating with the number of occupants. The system's output is the estimated number of people in the indoor environment. Hence, the downlink communication signal of the THz AP is used to sense the existing number of people on the indoor layout. The indoor people counting problem formulation, the utilized data preprocessing, and the proposed CNN-based indoor people counting model are discussed in the following subsections.

1) Problem Formulation: We consider a wireless THz network in a dynamic environment with mobile users in an indoor room layout. Given a fixed window of N historical wireless THz channel gain data points of the downlink signal, denoted as $\{g(t_1), g(t_2), ..., g(t_N)\}$ collected at discrete time instances $t_1, t_2, ..., t_N$, the task is to design a classification model that, given a new THz channel gain data points of the downlink signal, will correctly predict the class, i.e., number of people, to which the new data points belong.

2) Data Preprocessing: In this study, we employed data preprocessing to ensure the robustness and reliability of our proposed deep learning (DL) model. We utilized a sliding window approach to segment the THz channel gain time series data into overlapping windows. This technique enhances the DL model's ability to capture temporal dependencies and patterns within the data. Each window is treated as an individual sample for subsequent processing. Also, each sample was standardized to have a mean of zero and a standard deviation of one. This standardization process ensures that all features contribute equally to the model's learning process and helps accelerate the convergence during training.

3) 1D CNN Model Training: The dataset was divided into training and testing sets with a ratio of $80\colon 20$. This split ensures that the DL model is evaluated on a representative subset of the data it has not seen during training, providing an unbiased assessment of its performance. The training set was further divided into training and validation sets with a ratio of $80\colon 20$. The validation set is used to fine-tune the model and prevent overfitting, ensuring that the model generalizes well to unseen data. We shuffled the window samples across the dataset to eliminate any potential bias due to the order of the samples and to achieve better generalization. The dataset contains eight classes, each representing a unique number of users. Each sample was assigned a class based on the ground truth value. This class assignment is critical for the supervised learning process, allowing the model to learn the mapping between input features and the corresponding output classes. The class labels were one-hot encoded. This transformation converts the categorical class labels into a binary matrix representation, which DL algorithms require.

The hyperparameters of the proposed DL model are optimized using random search. The number of convolutional layers is dynamically selected between 1 to 5. The kernel size is dynamically selected between 2 to 5. The number of filters in each layer is dynamically selected between 64 to 512. The number of strides is dynamically selected between 1 to 3. The activation function is dynamically selected among ReLU, tanh, and sigmoid.

The proposed DL model is a 1D CNN consisting of five convolutional layers. The first layer consists of 512 filters, kernel size of three, one stride, and ReLU activation function. The second layer consists of 320 filters, kernel size of four, one stride, and tanh activation function. The third layer consists of 256 filters, kernel size of five, one stride, and tanh activation function. The fourth layer consists of 128 filters, kernel size of two, three strides, and tanh activation function. The fifth layer consists of 512 filters, kernel size of two, two strides, and tanh activation function. The output from the fifth convolutional layer is then fed into a global average pooling 1D layer, which reduces the spatial dimensions. Finally, a dense layer with eight neurons and a softmax activation function is employed to predict the model's output, representing the probabilities of the eight possible classes.

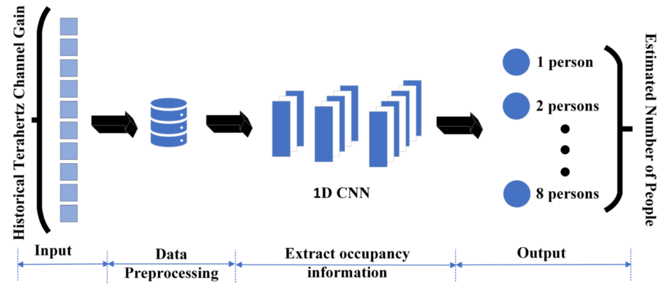

Fig. 4. The proposed indoor people counting system.

The DL model undergoes 200 epochs of training with a batch size of 32, using the Adam optimizer initialized at a learning rate of 0.001. During training over $272,000$ samples, 20% of the data is reserved for validation to monitor the model's generalization capability. The learning rate is reduced by 50% if no improvement in validation loss is observed for three consecutive epochs. Furthermore, the training process is terminated if there is no improvement in validation loss for six consecutive epochs, thereby preventing overfitting and conserving computational resources. This learning approach not only aims to achieve optimal model accuracy but also prevents overfitting and enhances the training efficiency by dynamically adjusting the learning rate and preventing unnecessary training epochs.

IV. Numerical Results

This section evaluates the proposed indoor people counting system. The frequency used during the THz dataset generation is 0.142 THz, which is included in the THz frequency band ranging from 0.1 to 10 THz [20]. Also, the window size N is equal to 60. Each model's training, validation, and testing were implemented based on the TensorFlow framework running with two AMD EPYC 9374F processors, each with 32 cores running at a base frequency of 3.85 GHz and 24 GB of RAM.

Fig. 5 shows the CNN model's loss of the training and validation datasets for the proposed indoor people counting system. From Fig. 5, the training loss steadily decreases and approaches near zero after 15 epochs. This shows that the model effectively minimizes errors in training data. Also, the validation loss follows a similar decreasing trend, although it is higher and fluctuates early on. It stabilizes around a low value after 30 epochs, indicating the model is not overfitting. The training and validation loss curves flatten out in the later epochs, indicating convergence. The small gap between the two losses after stabilization suggests good generalization. The model shows good convergence, with stable accuracy and low loss for training and validation datasets after around 30

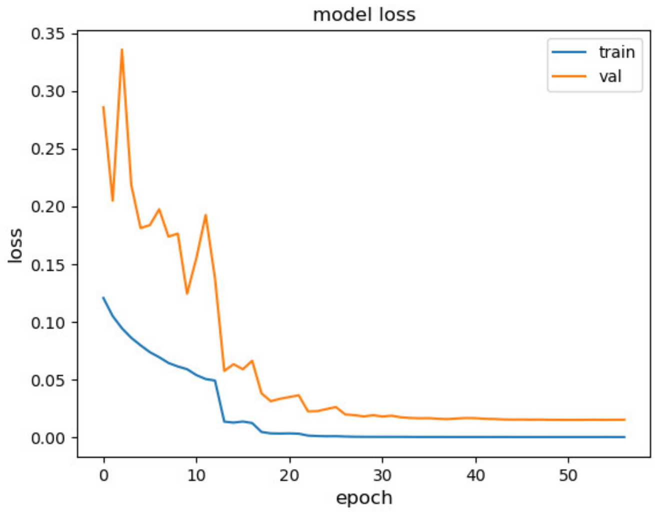

Fig. 5. The proposed CNN model's loss of the training and validation datasets for the proposed indoor people counting system.

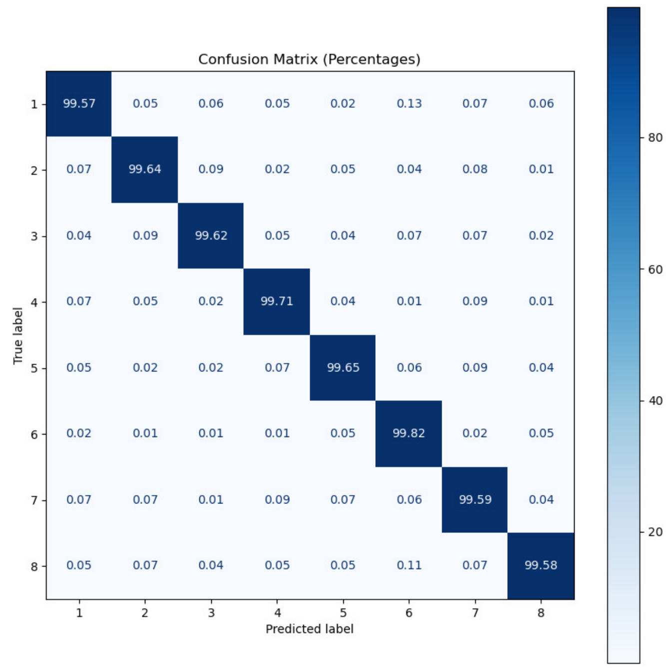

Fig. 6. The confusion matrix of the proposed indoor people counting system.

epochs. No significant gap between training and validation indicates that the model generalizes well without overfitting.

The confusion matrix in Fig. 6 demonstrates the CNN model's classification accuracy of the proposed indoor people counting system. Each class has been tested with over 8500 samples. The matrix entries are normalized and presented as percentages for more precise interpretation.

The high diagonal values in the confusion matrix indicate the proposed system's ability to estimate the number of people in the indoor environment accurately. The minimal off-diagonal values signify that misclassifications are rare. Overall, the results validate the effectiveness of the proposed indoor people counting system, with each class achieving near-perfect accuracy. With estimation accuracies exceeding 99.5% across all classes, the proposed system is highly reliable and can be effectively deployed in real-world scenarios for accurate indoor people counting.

The classification report demonstrated in Table I contains precision, recall, and F1-score metrics for eight classes, with each class tested against 8514 samples. The classification report shows that the proposed CNN model performs well across all classes, with high precision, recall, and F1-scores, reflecting a robust and reliable classifier in estimating the number of indoor people.

To better assess the proposed CNN classifier's performance in estimating the number of people indoors, we developed a random forest time series classifier to compare it with the proposed CNN classifier's performance. The random forest time series classifier has 100 individual decision trees, and each decision tree makes an independent prediction about the estimated number of people indoors. Then, the random forest output will be the majority estimated number of people predicted by the individual decision trees. The catch22 feature extractor proposed in [21] extracts the features that will be fed to the random forest classifier. In Table II, the performance is evaluated in terms of accuracy, precision, recall, and F1-score for estimating different numbers of people, ranging from one to eight. Table II indicates that the proposed CNN outperforms the random forest by more than 10% in terms of accuracy, precision, recall, and F1-score. However, the training time of the proposed CNN model is 23 times longer than that of the random forest model, as shown in Table III.

TABLE I
THE CLASSIFICATION REPORT OF THE PROPOSED CNN.

Class	Precision	Recall	F1-score	Support
1	0.9964	0.9957	0.9960	8514
2	0.9964	0.9964	0.9964	8514
3	0.9974	0.9962	0.9968	8514
4	0.9966	0.9971	0.9968	8514
5	0.9969	0.9965	0.9967	8514
6	0.9953	0.9982	0.9968	8514
7	0.9950	0.9959	0.9954	8514
8	0.9978	0.9958	0.9968	8514

TABLE II
THE PERFORMANCE OF THE PROPOSED CNN VS. RANDOM FOREST IN TERMS OF ACCURACY, PRECISION, RECALL, AND F1-SCORE.

Metric	CNN	Random Forest
Accuracy	0.9965	0.8889
Precision	0.9965	0.8894
Recall	0.9965	0.8889
F1-score	0.9965	0.8890

TABLE III
COMPLEXITY COMPARISON BETWEEN THE PROPOSED CNN AND THE RANDOM FOREST IN TERMS OF TRAINING TIME.

Model	**Training Time (s)**
CNN	16674
Random Forest	720

V. CONCLUSION

In this paper, we proposed an indoor people counting system based on the wireless THz communication signal. The proposed system is based on a 1D CNN time series classifier. The input to the proposed CNN time series classifier is the wireless THz channel gain of the THz AP downlink, whereas the output is the estimated number of people in the indoor layout. The numerical results demonstrate the effectiveness of the proposed system in counting up to eight people indoors with 99.5% accuracy. Also, we compared the performance of the proposed CNN time series classifier with a random forest time series classifier that uses catch22 feature extraction. The numerical results indicate that the proposed CNN model outperforms the random forest by more than 10% in terms of accuracy, precision, recall, and F1-score. However, the training time of the proposed CNN model is 23 times longer than that of the random forest model.

REFERENCES

[1] Y.-K. Cheng and R. Y. Chang, "Device-free indoor people counting using Wi-Fi channel state information for internet of things," in *GLOBECOM 2017 - 2017 IEEE Global Communications Conference*, 2017.

[2] M. E. Colmenar Cascón, "Cloud-based people counting system," Master's thesis, Universidad Pontificia Comillas, 2021.

[3] R. L. dos Santos, H. C. de Oliveira, M. C. de Almeida, D. F. Vieira, E. P. L. Junior, and T. Ji, "A low-cost bidirectional people counter device for assisting social distancing monitoring for COVID-19," *Journal of Control, Automation and Electrical Systems*, vol. 33, no. 4, pp. 1148–1160, 2022.

[4] R. Shahbazian and I. Trubitsyna, "Human sensing by using radio frequency signals: A survey on occupancy and activity detection," *IEEE Access*, vol. 11, pp. 40 878–40 904, 2023.

[5] H. Jiang, S. Chen, Z. Xiao, J. Hu, J. Liu, and S. Dustdar, "Pa-Count: Passenger counting in vehicles using Wi-Fi signals," *IEEE Transactions on Mobile Computing*, vol. 23, no. 4, pp. 2684–2697, 2024.

[6] N. Krishnachaithanya, G. Singh, S. Sharma, R. Dinesh, S. R. Sihag, K. Solanki *et al.*, "People counting in public spaces using deep learning-based object detection and tracking techniques," in *2023 International Conference on Computational Intelligence and Sustainable Engineering Solutions (CISES)*, 2023, pp. 784–788.

[7] X. Yang, J. Gaspar, W. H. Lou, W. Ke, C. T. Lam, and Y. Wang, "Vision-based mobile people counting system," in *Proceedings of the 2019 4th International Conference on Machine Learning Technologies*, ser. ICMLT '19, 2019, p. 42–46.

[8] P.-R. Tsou, C.-E. Wu, Y.-R. Chen, Y.-T. Ho, J.-K. Chang, and H.-P. Tsai, "Counting people by using convolutional neural network and a PIR array," in *2020 21st IEEE International Conference on Mobile Data Management (MDM)*, 2020, pp. 342–347.

[9] S. T. Kouyoumdjieva, P. Danielis, and G. Karlsson, "Survey of non-image-based approaches for counting people," *IEEE Communications Surveys & Tutorials*, vol. 22, no. 2, pp. 1305–1336, 2020.

[10] J. Yang, M. Yang, Y. Wu, P. Yang, and Z. Liu, "Wi-PC: A device-free people counting system based on WiFi channel state information," in *2023 IEEE Smart World Congress (SWC)*, 2023.

[11] J. R. M. Bernaola, I. Sobrón, J. Del Ser, I. Landa, I. Eizmendi, and M. Vélez, "Ensemble learning for seated people counting using WiFi signals: Performance study and transferability assessment," in *2021 IEEE Globecom Workshops (GC Wkshps)*, 2021.

[12] J.-H. Choi, J.-E. Kim, and K.-T. Kim, "Deep learning approach for radar-based people counting," *IEEE Internet of Things Journal*, vol. 9, no. 10, pp. 7715–7730, 2022.

[13] C.-C. Shih, X. Zhou, T. Nguyen, and K. Pham, "People counting system using mmwave MIMO radar with 3D convolutional neural network," in *2023 IEEE 97th Vehicular Technology Conference (VTC2023-Spring)*, 2023.

[14] L. Ren, A. G. Yarovoy, and F. Fioranelli, "Grouped people counting using mm-Wave FMCW MIMO radar," *IEEE Internet of Things Journal*, vol. 10, no. 22, pp. 20 107–20 119, 2023.

[15] G. Ma, S. Dang, M.-S. Alouini, and B. Shihada, "Smart buildings enabled by 6G communications," *IEEE Internet of Things Magazine*, vol. 5, no. 2, pp. 181–186, 2022.

[16] Z.-Y. Wu, M. Ismail, J. Kong, E. Serpedin, and J. Wang, "Channel characterization and realization of mobile optical wireless communications," *IEEE Transactions on Communications*, vol. 68, no. 10, pp. 6426–6439, 2020.

[17] E. Hasan, E. Mahalal, M. Ismail, Z.-Y. Wu, M. M. Fouda, T. Koketsu Rodrigues, and N. Kato, "Robust deep learning-based indoor mmWave channel prediction under concept drift," in *2023 IEEE 98th Vehicular Technology Conference (VTC2023-Fall)*, 2023.

[18] F. Sheikh, Y. Gao, and T. Kaiser, "A study of diffuse scattering in massive MIMO channels at Terahertz frequencies," *IEEE Transactions on Antennas and Propagation*, vol. 68, no. 2, pp. 997–1008, 2020.

[19] C. Han, A. O. Bicen, and I. F. Akyildiz, "Multi-ray channel modeling and wideband characterization for wireless communications in the Terahertz band," *IEEE Transactions on Wireless Communications*, vol. 14, no. 5, pp. 2402–2412, 2015.

[20] W. Jiang, Q. Zhou, J. He, M. A. Habibi, S. Melnyk *et al.*, "Terahertz communications and sensing for 6G and beyond: A comprehensive review," *IEEE Communications Surveys & Tutorials*, early access, doi: 10.1109/COMST.2024.3385908.

[21] C. H. Lubba, S. S. Sethi, P. Knaute, S. R. Schultz, B. D. Fulcher, and N. S. Jones, "catch22: Canonical time-series characteristics: Selected through highly comparative time-series analysis," *Data Mining and Knowledge Discovery*, vol. 33, no. 6, pp. 1821–1852, 2019.