









# A Feasibility Study on Indoor Localization and Multiperson Tracking Using Sparsely Distributed Camera Network With Edge Computing

Hyeokhyen Kwon , Chaitra Hegde , Yashar Kiarashi, Venkata Siva Krishna Madala , Ratan Singh, ArjunSinh Nakum , Robert Tweedy , Leandro Miletto Tonetto , Craig M. Zimring, Matthew Doiron, Amy D. Rodriguez , Allan I. Levey, and Gari D. Clifford , *Fellow, IEEE*

**Abstract**—Camera-based activity monitoring systems are becoming an attractive solution for smart building applications with the advances in computer vision and edge computing technologies. In this article, we present a feasibility study and systematic analysis of a camera-based indoor localization and multiperson tracking system implemented on edge computing devices within a large indoor space. To this end, we deployed an end-to-end edge computing pipeline that utilizes multiple cameras to achieve localization, body orientation estimation, and tracking of multiple individuals within a large therapeutic space spanning 1700 m<sup>2</sup>, all while maintaining a strong focus on preserving privacy. Our pipeline consists of 39 edge computing camera systems equipped with tensor processing units (TPUs) placed in the indoor space's ceiling. To ensure the privacy of individuals, a real-time multiperson pose estimation algorithm runs on the TPU of the computing camera system. This algorithm extracts poses and bounding boxes, which are utilized for indoor localization, body orientation estimation, and multiperson tracking. Our pipeline demonstrated an average localization error of 1.41 m, a multiple-object tracking accuracy score of 88.6%, and a mean absolute body orientation error of 29°. These results show that localization and tracking of individuals in a large indoor space is feasible even with the privacy constrains.

**Index Terms**—Body orientation estimation, cloud computing, computer vision, edge computing, indoor localization, multiperson tracking.

## I. INTRODUCTION

INDOOR localization and tracking is an active research area to support various applications, such as smart hospitals [1],

Manuscript received 29 June 2023; revised 13 November 2023; accepted 22 November 2023. Date of publication 28 November 2023; date of current version 18 December 2023. This work was supported by James M. Cox Foundation and Cox Enterprises, Inc., through Emory's Brain Health Center and Georgia Institute of Technology. (Hyeokhyen Kwon and Chaitra Hegde contributed equally to this work.) (Corresponding author: Hyeokhyen Kwon.)

Hyeokhyen Kwon, Yashar Kiarashi, and Robert Tweedy are with the Department of Biomedical Informatics, School of Medicine, Emory University, Atlanta, GA 30322 USA (e-mail: hyeokhyen.kwon@emory.edu; yash@dbmi.emory.edu; robert.tweedy@emory.edu).

Chaitra Hegde, Venkata Siva Krishna Madala, Ratan Singh, and ArjunSinh Nakum are with the School of Electrical and Computer Engineering, Georgia Institute of Technology, Atlanta, GA 30332 USA (e-mail: chegde@gatech.edu; vmadala3@gatech.edu; rsingh388@gatech.edu; arjun5inh@gatech.edu).

Leandro Miletto Tonetto is with the School of Industrial Design, College of Design, Georgia Institute of Technology, Atlanta, GA 30332 USA (e-mail: ltonetto3@gatech.edu).

Craig M. Zimring is with the School of Architecture, College of Design, Georgia Institute of Technology, Atlanta, GA 30332 USA (e-mail: craig.zimring@design.gatech.edu).

Matthew Doiron, Amy D. Rodriguez, and Allan I. Levey are with the Department of Neurology, School of Medicine, Emory University, Atlanta, GA 30322 USA (e-mail: matthew.james.doiron@emory.edu; amy.rodriguez@emory.edu; alevey@emory.edu).

Gari D. Clifford is with the Department of Biomedical Informatics, Emory University School of Medicine, Atlanta, GA 30322 USA, also with the Department of Biomedical Engineering, Georgia Institute of Technology, Atlanta, GA 30332 USA, and also with the Emory University, Atlanta, GA 30322 USA (e-mail: gari@gatech.edu).

Code available at <https://github.com/cliffordlab/EdgeCameraTracking>  
Digital Object Identifier 10.1109/JISPIN.2023.3337189

[2], smart offices [3], or customer analysis in shopping mall [4]. A wide range of sensors have been used to simultaneously detect the whereabouts and movement of multiple people in indoor spaces over the past two decades, including radio-frequency identification [5], infrared [6], WiFi [7], and Bluetooth [8]. Recent advances in multiperson pose estimation [9], [10] and tracking have led to the use of cameras to monitor indoor activities [11], [12], and vision-based approaches have proven to be very effective in capturing accurate and detailed movements continuously over time. However, most studies have only examined a few cameras in small rooms and have not studied a system that integrates localization, tracking, and body orientation estimation techniques in a large indoor space required to understand fine-grained activities comprehensively [1], [13], [14]. Advancements in edge computing now allow for low-cost, powerful devices to monitor activity in existing buildings [15], [16], [17].

In this work, we present a feasibility study and systematic analysis of a camera-based indoor localization and tracking system to monitor multiple people across a wide indoor space using an edge computing framework. To carry out this study, we deployed a multicamera edge computing pipeline designed to estimate locations, track individuals, and ascertain their body orientations in a vast therapeutic indoor space with a strong emphasis on safeguarding privacy. We deployed 39 cameras sparsely positioned throughout a large indoor space spanning 1700 m<sup>2</sup>, which contains diverse functional areas, such as a gym, library, and kitchen. These cameras were connected to edge computing units, specifically Raspberry Pi v4 B devices. We adopted a real-time processing approach using Google Coral

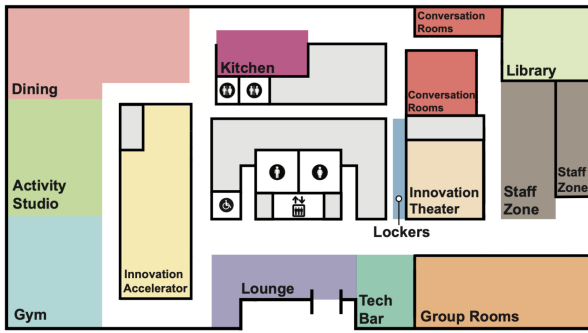


Fig. 1. Space usage in the study space.

tensor processing units (TPUs) connected to the edge computing units. These TPUs automatically process raw color frames obtained from the cameras to detect 2-D poses (as described in [18]). This information is then used to estimate the positions and orientations of people within the indoor space, and the original raw frames are discarded. The estimated positions and orientations from all the cameras are combined to track the movement of individuals. Using our camera infrastructure, we aim to provide a comprehensive analysis of the camera installation and its relations with the performance of our system in various activity scenarios. We expect this work to provide practical guidelines for practitioners when installing a multiview camera-based activity monitoring system using edge computing frameworks in existing facilities.

## II. METHOD AND MATERIALS

### A. Study Site

Our site was designed to promote the involvement of older adults in a range of therapeutic activities, including physical exercise and memory training. It also provides opportunities for social interaction, such as shared meals and comfortable lounge areas. The study site consists of diverse functional regions, as shown in Fig. 1. The activity area (A) encompasses both the activity studio and gym, promoting physical exercise. The kitchen (B) covers the upper corridor between the dining area's right boundary and the conversation room's left boundary, encouraging independent cooking activities. The left corridor (C) runs alongside a room called the innovation accelerator, whereas the right corridor (E) is situated next to the lockers, to the left of the innovation theater. The lounge (D) extends from the elevator area to the front door, and the staff area (F) comprises the section on the building's right side, where staff offices are located.

### B. Distributed Camera Network Using Edge Computing

1) *Camera Configuration*: Fig. 2(a) shows the locations (blue) of 39 low-cost (< \$500) edge computing and camera devices [see Fig. 2(b)] and their viewpoints (red). These cameras effectively provide coverage over the entire study site where subjects have consented to have their activities monitored. The cameras along with the edge devices were placed in the ceiling of the indoor environment away from the direct line of sight

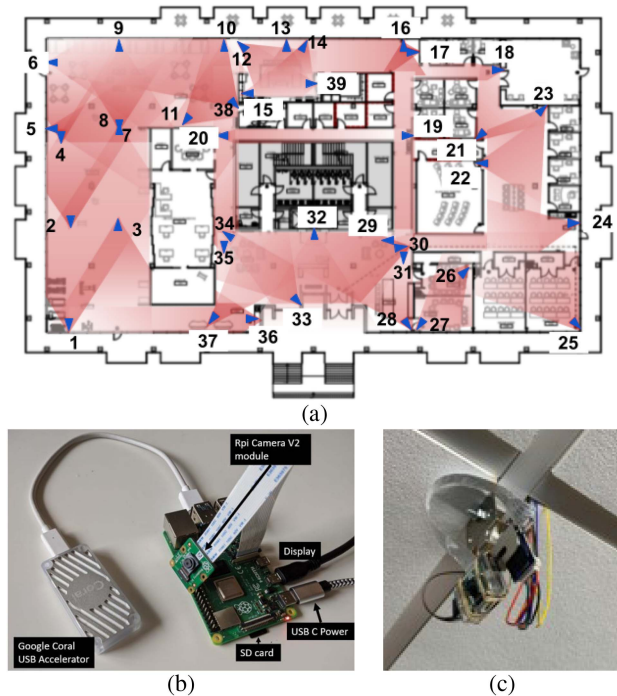


Fig. 2. Cameras installation in the study site. (a) In total, 39 camera sensors (Sony IMX219 8-megapixel sensor; blue) are installed with the field-of-view (red) covering the entire indoor space. (b) Raspberry Pi 4 is installed with Google Coral TPU USB Accelerator to run deep learning-based 2-D pose estimation models in real time. (c) Example of camera installation in the ceiling. Raspberry Pi 4 is hidden in the ceiling and connected to the nearest network and power sources.

of visitors to prevent them from feeling overly self-conscious, as shown in Fig. 2(c). These edge devices were connected to the closest available power and network sources in the ceiling through universal serial bus type-C (USB-C) power cable and Ethernet cable for data transfer. Depending on the availability of power and network sources, some locations, such as kitchen or staff areas, are covered by multiple cameras, and other locations, such as long corridors, are covered by a single or two cameras. The cameras installed above visitors' eyesight make it easier for visitors to comply with the proposed instrumentation of cameras for activity monitoring.

2) *Secure Computing Network Architecture*: The complete computing setup comprises four key components: the offsite cloud server, the onsite fog server, the edge computing network infrastructure, and the array of 39 edge computing camera devices. Here, onsite and offsite servers refer to servers that are physically present in the study site and servers places outside of the study site, respectively. This configuration is shown in Fig. 3.

Real-time estimation of 2-D poses, composed of 17 keypoints following the MS COCO format [19], and bounding boxes occurs at a frame rate of 1 Hz using the TPU on the edge computing device. Bounding box images are temporarily retained on the onsite fog computing node for up to one hour to facilitate body orientation estimation. They are then discarded. The raw frames are immediately discarded from the edge computing device after the 2-D poses are extracted and do not get sent to the onsite




Fig. 3. Computing network architecture and data management in our study.

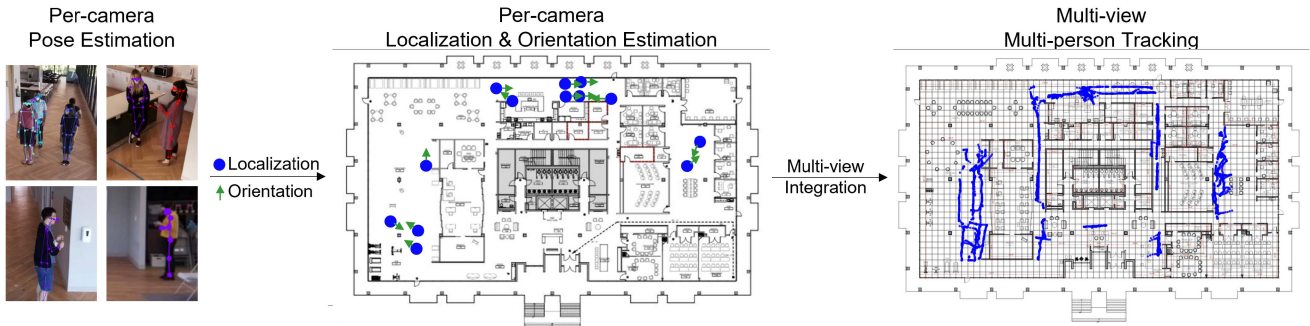


Fig. 4. Overall diagram of the proposed activity monitoring system. From each camera, the location and orientation of multiperson poses are projected into the study space's 2-D coordinates. While visitors spend time in the study space, samples from multiple views are integrated and tracked.

fog server. This ensures privacy of individuals in the study site and prevents the capture, storage, or transmission of sensitive or identifiable information, such as faces or nudity. The 2-D poses and body orientations of a single day, specifically from 8 A.M. to 7 P.M., are stored on the fog server and these data are routinely synchronized with the offsite cloud server during the night, typically between 12 A.M. and 7 A.M. The duration of the nightly synchronization can change depending on the volume of data collected during the day. It is worth noting that the fog server complies with the Health Insurance Portability and Accountability Act regulations.

### C. Monitoring Movement in Built-in Environment

1) *Overall Pipeline*: The overall approach for multiview multiperson tracking using the abovementioned computing and camera infrastructure is shown in Fig. 4. Using the 2-D poses and human bounding boxes, each person's location and body orientation is estimated and projected on a 2-D coordinate system representing the study space. Multiview integration then resolves duplicate detections across cameras from overlapping views by combining samples from 39 cameras. Lastly, a multiperson tracking algorithm [20] is used to track people from the combined location and orientation samples.

2) *Preprocessing 2-D Poses*: The 2-D poses estimated by the edge computing systems undergo a preprocessing step to eliminate noisy data. We identified two primary issues in the poses detected by the pose detection model [18]: nonhuman poses and temporal inconsistencies. The model occasionally misinterprets patterns on the walls and floor as human poses due to its overhead camera perspective. These nonhuman poses remain static over time and show minimal pose changes.

Consequently, we identify and remove these false positives (nonhuman samples) by examining subsequent frames for nearzero pose changes. The second issue relates to the lack of temporal consistency in 2-D poses because the model estimates poses frame-by-frame. Some poses exhibit discrepancies, such as left and right sides being flipped or even missing entirely between frames (false negatives). To address this, we implement temporal smoothing using a Kalman filter-based approach to reduce noise and fill in missing pose detections [21], [22].

3) *Indoor Localization*: Our indoor localization technique is designed to determine the positions of individuals within the study site using a 2-D coordinate system that represents the layout of the site, which we call the study site coordinate system. This is achieved by utilizing the 2-D poses detected from each camera. In our approach, we define the location of a person within a frame as the midpoint between their left and right feet. Given that the floor in our study area is uniformly flat, all individuals detected in a frame are assumed to be on the same plane. To enable this localization, we derive a perspective transformation matrix  $M \in \mathbb{R}^{3 \times 3}$ . This matrix serves to project the foot locations from the 2-D camera space to the corresponding floor positions within the study site coordinate system [12], [23]. The process for obtaining this transformation matrix  $M$  is carried out manually for each camera by aligning four corresponding points in both the 2-D camera space and the study space coordinate system.

4) *Body Orientation Estimation*: Starting with the human bounding boxes obtained from 2-D pose detection models, we extract the chest-facing direction of each individual within the transverse plane. This orientation is subsequently translated into the study space coordinate system. To estimate the chest-facing direction from a human bounding box image, we employ a deep



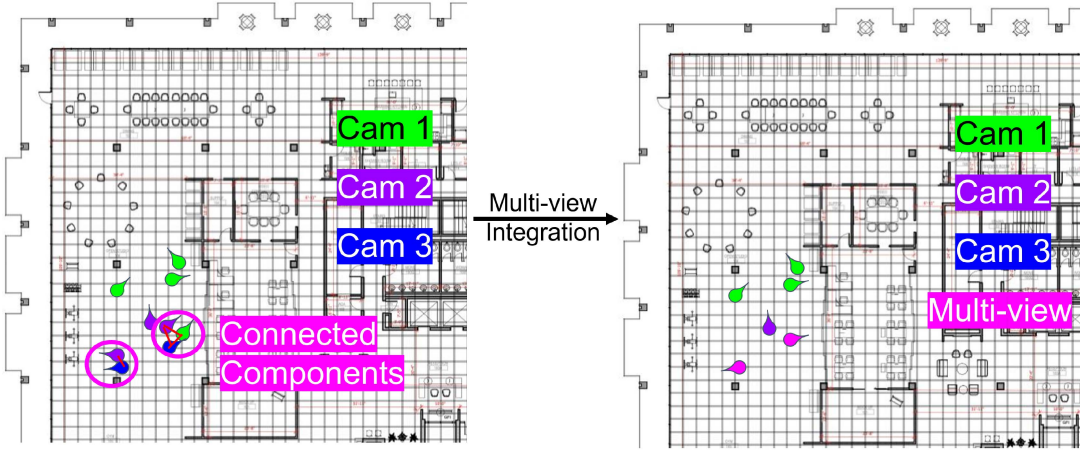


Fig. 5. Multiview integration of duplicated samples from multiple cameras having overlapping views. We first generate a graph, in which vertices are all samples from all cameras. Then, the edges are created for samples from different cameras that are close together (within 1.5 m) (left-hand side, red). The connected components are then integrated as a single sample potentially representing the true location and orientation of the person in the study space (right-hand side, magenta).

learning-based method for body orientation estimation [24]. Here, we define a human bounding box as a tightly enclosed square region that encompasses a detected 2-D pose from the edge computing systems. The resultant estimate of the chest orientation is represented as a 2-D chest vector on the transverse plane. This 2-D chest vector, initially defined in camera space, is then transformed into a 2-D orientation in the study site coordinate system, denoted as  $v_{\text{world}}^t = R_{\text{camera} \rightarrow \text{world}} \vec{v}_{\text{world}}^t$ . The rotation matrices,  $R_{\text{camera} \rightarrow \text{world}}$ , are manually determined for each camera based on the direction in which the camera is facing within the study site, with north serving as the reference orientation.

5) *Multiview Integration*: To eliminate duplicated samples that may occur due to overlapping camera fields-of-view, we adopt a method that involves detecting people across multiple cameras and consolidating these detections into a single sample, potentially representing the person's actual location, as illustrated in Fig. 5. To identify duplicate detections, we construct a graph, denoted as  $G = (V, E)$ . In this graph, the vertices, represented as  $V = \{v_{ci} = (x, y) | c = 1, \dots, 39, i = 1, \dots\}$ , correspond to the locations of samples from all 39 cameras. The edges, denoted as  $E = \{v_{ci}v_{c'j} | c \neq c'\}$ , signify the connections between samples from different cameras that are sufficiently close in distance. We assume that duplicates originate from the same person captured by different cameras if their Euclidean distance is within  $\|v_{ci} - v_{c'j}\|_2 \leq 1.5$  m. This threshold takes into account the potential errors in our localization technique. While variations may occur due to lens distortions, which can lead to discrepancies in the perspective transformations, our empirical findings suggest that a margin of 1.5 m is adequate [25].

Once the graph  $G$  is established, duplicate samples are identified as connected components (CC) within the graph, following principles from graph theory [26]. The integration of duplicate sample locations considers the distance of cameras from the detected sample. More precisely, during the integration of duplicated samples, a penalty is applied when the distance between the estimated person's location and the camera is substantial. This approach acknowledges the strong correlation between

localization errors and the distance between the camera and the target person in our analysis. This is represented by the following:

$$L = \sum_{\text{CC}} \frac{1}{d_{ci}} v_{ci}. \quad (1)$$

Here,  $L$  denotes the integrated sample location, and  $d_{ci} = \|l_c - v_{ci}\|_2^2$  signifies the Euclidean distance between the sample,  $v_{ci}$ , and the camera,  $l_c$ .

Similarly, the orientation of the duplicated samples is integrated for the vertices within the CC, taking camera distances into account. This multiview integration process results in unique locations, denoted as  $L = [l_x, l_y]$ , and orientation vectors,  $o = [o_x, o_y]$  for each person at each time step.

6) *Multiperson Tracking*: We implement Kalman filter-based tracking for the locations denoted as  $L = [l_x, l_y]$  and the orientation vector represented by  $O = [o_x, o_y]$ . This approach assumes that people's facing orientation closely follows their walking direction [27] and can be expressed as follows:

$$O_t = O_{t-1} + w \cdot \dot{L}_{t-1} + dt \cdot \dot{O}_{t-1}. \quad (2)$$

Here, the parameter  $w = 0.1$  acts as a weighting factor for the influence of the movement direction (or velocity) on the current orientation, and  $dt$  denotes the time interval between frames, which is set at 1 s. To further enhance tracking performance, we apply Rauch–Tung–Striebel smoothing to completed tracks, which takes into account temporal dynamics across the entire trajectory of a person.

#### D. Benchmark Data Collection

In order to methodically assess the effectiveness of our system for localization, body orientation estimation, and tracking, we created an annotated dataset. This dataset features four to five individuals moving and interacting within the study site. The benchmark dataset encompasses scenarios where people are 1) walking together, 2) walking past each other, or 3) conversing in place while adjusting their facing orientations. We marked the

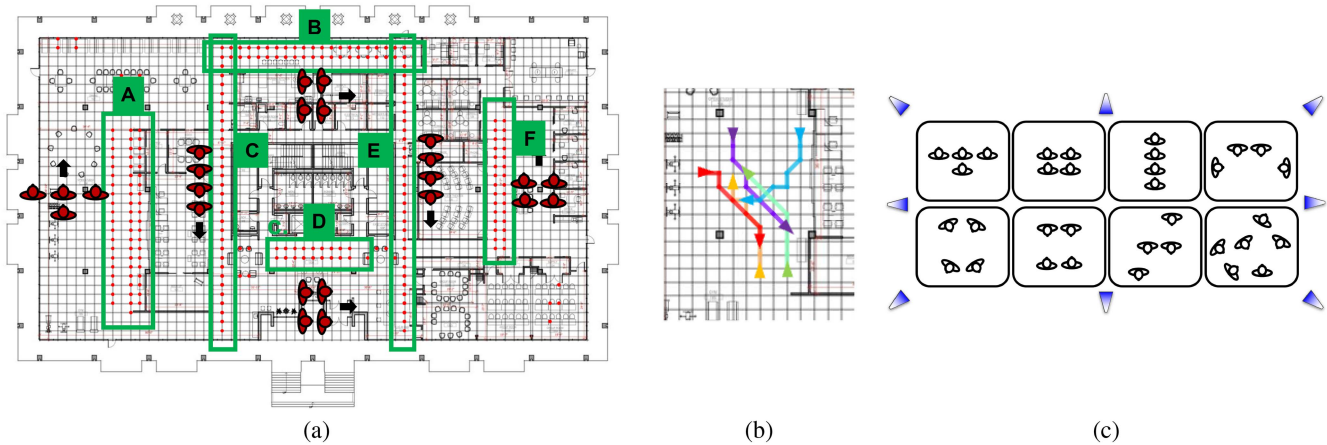


Fig. 6. Benchmark data collection. We selected the group formations in order to simulate realistic scenarios during the study. (a) Four participants (red) were walking with different formations while collecting the benchmark dataset. Our team has covered nearly all areas of the study site where visitors spend a majority of their time. (b) Five participants showing more complex movements with occlusions. This was collected in activity area, shown in (a) A, where most physical activity occurs in our study site. Each person is color-coded with different colors. (c) Overall examples of group formations taken by participants and camera viewpoints while collecting the benchmark data.

entire area with 1-m intervals to provide ground-truth reference points. The subjects were instructed to follow predetermined group formations and designated walking paths. As the individuals moved within the space, an observer recorded the time and location of each subject at 1-s intervals. In addition, the observer manually annotated the body orientation of the subjects, using the north direction (as indicated by the upper side of Fig. 1) as a reference point. Below, we outline the three distinct datasets we generated to capture various activities within the space.

1) *Four People Walking Around Entire Space*: In the first benchmark dataset, we had four individuals moving throughout the study site. To ensure that our dataset encompasses realistic use cases for visitors in the study site, we carefully selected locations and paths, as depicted in Fig. 6. These areas were of particular interest for tracking behaviors within the study space [see Fig. 6(a)]. These locations included spaces frequently used by visitors for activities, such as dining, kitchen usage, physical exercises, relaxation, and staff zones (see Fig. 1), where participants had given their consent to be recorded.

For annotation purposes, we utilized the closest 1-m markers as the ground-truth reference points for the subjects while they were walking. This level of detail was sufficient to account for variations in foot size and stride length among the participants. As the subjects walked, we introduced changes in their group formation to simulate realistic movement patterns within the space [see Fig. 6(c)]. In wider corridors, we arranged two or three people to walk side-by-side, whereas in narrower corridors, multiple subjects proceeded in straight lines. This approach allowed us to assess the impact of occlusions, where subjects might not be visible from certain camera viewpoints.

2) *Five People Walking Past Each Other*: As depicted in Fig. 6(b), we recreated a scenario in the activity area where multiple individuals walked toward one another and passed by each other, a situation that often results in occlusions. Such occlusions are common during activities, such as dancing or group exercises.

3) *Five People Changing Orientations in Place*: During group conversations, individuals often remain seated or stand

stationary, adjusting their facing directions to interact. In this dataset, we focused on situations where people made subtle changes in their facing orientation while staying in one spot, as shown in Fig. 7. We conducted this dataset in three sessions, starting with a smaller social group of two participants and gradually increasing it to five. In this context, a social group refers to people engaging in active communication within a distance of less than 2 m, taking into account pandemic-related social distancing guidelines. Participants followed instructions to gradually modify their facing direction as part of the data collection process.

### E. Evaluating Multiperson Tracking and Localization

We assess the performance of our activity monitoring system by examining its localization, tracking, and body orientation estimation capabilities, following established evaluation methods [24], [28]. Since there is not a singular metric for evaluating multiperson tracking methods, we employ a set of metrics to comprehensively evaluate the proposed system [29]. In the evaluation, metrics denoted by  $\uparrow$  indicate better performance when higher, whereas those marked with  $\downarrow$  signify better performance when lower.

- 1) Multi-object tracking accuracy (MOTA)  $\uparrow$  combines three kinds of error sources, which are false positives, missed targets, and identity switches.
- 2) ID F1 score (IDF1)  $\uparrow$  is the ratio of correctly identified detections over the average number of ground-truth and computed detections.
- 3) Mostly tracked (MT) targets  $\uparrow$  measures the ratio of ground-truth trajectories that are covered by a track hypothesis for at least 80% of their respective life span.
- 4) Mostly lost (ML) targets  $\downarrow$  measures the ratio of ground-truth trajectories that are covered by a track hypothesis for at most 20% of their respective life span.
- 5) False positives rate (FPR)  $\downarrow$  measures the average number of false alarms per frame.

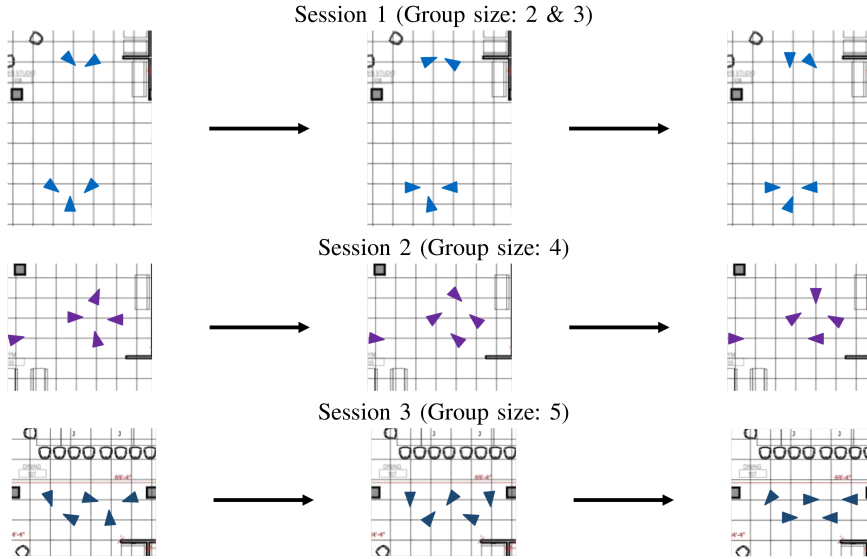


Fig. 7. Data collection when people make orientation changes while standing. We capture three sessions from different locations in the activity area, with different sizes of social groups. The first session had two social groups with two and three members, the second session had one social group with four members, and the third session had one social group with five members. Participants were instructed to change their facing directions during conversations.

- 6) False negative rates (FNRs)  $\downarrow$  measures the average number of missed targets per frame.
- 7) Recall (Rcll)  $\uparrow$  measures the ratio of correct detections to the total number of ground-truth locations.
- 8) Precision (Prcn)  $\uparrow$  measures the number of detected objects over the sum of detected and false positives.
- 9) ID switch ratio (IDSR)  $\downarrow$  measures the number of ID switches (IDS) over Rcll.
- 10) Fragmentation (Frag)  $\downarrow$  measures the total number of times a trajectory is fragmented, which are interruptions during tracking.

In particular, we assess the system's performance using the MOTA metric, defined as follows:

$$\text{MOTA} = 1 - \frac{\sum_t \text{FN}_t + \text{FP}_t + \text{IDS}_t}{\sum_t \text{GT}_t}. \quad (3)$$

Here,  $\text{FN}_t$ ,  $\text{FP}_t$ ,  $\text{IDS}_t$ , and  $\text{GT}_t$  represent false negatives (missed detections), false positives, IDS, and ground truths at time  $t$ , respectively.

An essential evaluation that precedes tracking is detection. Here, we define a ground-truth target at time  $t$  as "detected" if the estimated sample location falls within a 1.5-m Euclidean distance. Otherwise, the target is considered "missed." The choice of a 1.5-m threshold is based on the observation that people typically walk at speeds below 1.42 m/s in everyday situations [30], [31], [32]. To showcase the accuracy of the estimated sample locations, we also report the multiple object tracking precision (MOTP)

$$\text{MOTP} = \frac{\sum_{i,t} d_t^i}{\sum_t c_t}. \quad (4)$$

Here,  $d_t^i$  represents the Euclidean distance between the  $i$ th detected ground-truth sample and the estimated sample, and  $c_t$

is the number of detections at time  $t$ . In simpler terms, MOTP provides the average Euclidean distance error for the detected samples, indicating the system's ability to estimate the precise location of target samples, independent of tracking trajectories or IDS.

#### F. Evaluating Body Orientation Estimation

Consistent with previous definitions of body orientation found in earlier works [24], [33], [34], we represent human orientation, denoted as  $\theta$ , within the range of  $[0^\circ, 360^\circ)$ . This angle corresponds to the direction in which the chest is facing in the study site coordinate system. The north (top) direction serves as the reference point, with clockwise rotation applied. For the assessment of body orientation, we utilize two metrics as also employed in previous studies [24], given in the following.

- 1) *Mean absolute error (MAE) of angular distance*: This metric quantifies the average absolute angular difference between predicted and ground-truth orientations.
- 2) *Accuracy- $X^\circ$* : Accuracy- $X^\circ$  is defined as the proportion of predictions falling within a deviation of  $X^\circ$  compared with the ground-truth orientation. By considering different values of  $X$  within the range of  $[0^\circ, 180^\circ]$ , this metric provides a fine-grained measure of model accuracy. Specifically, we evaluate accuracy- $X^\circ$  for a set of  $X$  values including  $X \in \{5^\circ, 15^\circ, 22.5^\circ, 30^\circ, 45^\circ, 90^\circ\}$  [24].

#### G. Impact of Camera Installation

We conducted a quantitative analysis to delve deeper into the influence of camera placement on four distinct factors that impact the performance of our system, as illustrated in Fig. 8. In particular, we aim to establish the correlation between the subjects' positions and facing orientations as captured within the video frames of each camera, and how these aspects affect






Fig. 8. Analyzing the impact of camera installation with respect to four aspects. (a) Distance between the camera and the person in the space. (b) Person's facing direction away from the camera. (c) Horizontal angle of the captured person from the camera center. (d) Vertical angle of the captured person's feet location from camera center. The horizontal angle (c;  $\theta_H^\circ$ ) and vertical angle (d;  $\theta_V^\circ$ ) are normalized by horizontal ( $\theta_{HFOV}^\circ$ ) and vertical ( $\theta_{VFOV}^\circ$ ) field-of-view angles, respectively, referring to the camera datasheet (Sony IMX219 8-megapixel sensor).

the accuracy of both localization and body orientation estimation. This analytical approach sheds light on the error budget arising from per-camera indoor localization (see Section II-C3) and body orientation estimation (see Section II-C4), ultimately contributing to our understanding of the overall system performance.

1) *Distance Between Camera and Person*: For all cameras, we assessed the distance between individuals captured in relation to the camera's center, as depicted in Fig. 8(a). This distance measurement is instrumental in evaluating the impact on localization and orientation errors.

2) *Facing Angle With Respect to Camera Viewpoint*: We gathered data on the extent to which the facing directions of individuals deviate from the camera viewpoints, as visualized in Fig. 8(b). Subsequently, we performed a correlation analysis to assess the relationship between these deviations and the associated localization and orientation errors.

3) *Horizontal Displacement From Camera Center*: We studied the impact of the horizontal displacement of individuals from the camera center, as illustrated in Fig. 8(c). To quantify this horizontal displacement, we introduced a normalized horizontal angle, denoted as  $\theta_H^\circ / \theta_{HFOV}^\circ$ , where  $\theta_H^\circ$  represents the horizontal angle of the person from the camera center, and  $\theta_{HFOV}^\circ$  is the horizontal field-of-view.

When a person moves closer to the horizontal center of the camera, the normalized horizontal angle approaches zero. Conversely, it increases to a value of 1 when the person reaches the edge of the frame. If the value exceeds 1, it indicates that only a portion of the person's body is visible within the camera's field-of-view. We explore how localization and orientation errors are influenced by the degree of horizontal displacement.

4) *Vertical Displacement From Camera Center*: We investigated the impact of vertical displacement of an individual from the camera center, as indicated in Fig. 8(d). To measure this vertical displacement, we introduced a normalized vertical angle denoted as  $\theta_V^\circ / \theta_{VFOV}^\circ$ . Here,  $\theta_V^\circ$  represents the vertical angle at which the person is captured from the camera center, and  $\theta_{VFOV}^\circ$  is the vertical field-of-view, which is obtained from the camera datasheet. We specifically focused on the foot location

as a reference point to calculate the normalized vertical angle. When a person is closer to the camera, the normalized vertical angle increases, and only the upper body is partially visible in the camera frame due to the limited vertical field-of-view. Conversely, as the person moves away from the camera, the normalized vertical angle decreases, allowing for the full body to be captured within the field-of-view. However, this comes at the cost of increased distance between the camera and the person. We examined how changes in localization and orientation errors are influenced by the vertical displacement of the captured individual.

### III. RESULTS

#### A. Multiperson Tracking and Localization

1) *Four People Walking Around Entire Space*: In this section, we present the system's overall performance when four people walk together throughout the entire study site. The results of multiperson tracking and localization are depicted in Fig. 9, and detailed in Table I, where each functional area is assessed individually. It is worth noting that different locations within the study site have varying sizes of the captured area (as shown in the second column of Table I) and a different number of cameras (as indicated in the third column of Table I), all of which captured 85% of the samples. The overall MOTA score stands at 62.9%, and the localization performance, measured by MOTP, is 1.41 m. These results indicate that our system achieved accurate localization and tracking in over 60% of cases, taking into account false positives, false negatives, and identity switches. A more detailed analysis of the impact of camera installation on localization accuracy is provided in Section III-C.

2) *Five People Walking Past Each Other*: Both qualitative and quantitative outcomes are presented in Fig. 9(b) and in the last row of Table I, respectively. Despite the challenging scenarios, the system demonstrated robust performance when individuals were passing by each other. The MOTA score, MOTP score, and IDF1 reached 79.8%, 1.44 m, and 70%, respectively, indicating strong performance with minimal ID switch errors in situations where people were crossing paths.

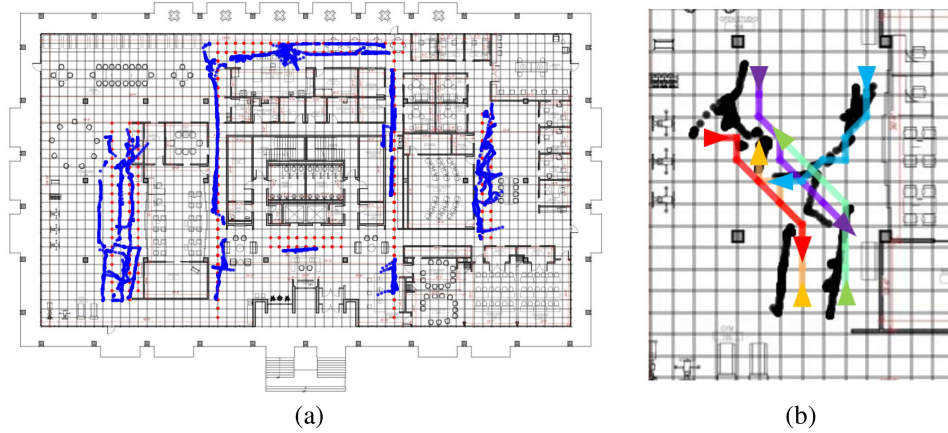


Fig. 9. Qualitative results of multiperson tracking. (a) Four people walking around the study site following Fig. 6(a). The reds are ground truth and blues are predicted location of people. (b) Five people walking in a complex pattern. Ground-truth trajectories are shown in colored lines with arrows indicating walking directions, whereas predicted trajectories are shown in black lines.

TABLE I  
MULTI-OBJECT LOCALIZATION AND TRACKING PERFORMANCE AT EACH REGION IN THE STUDY SITE WHEN FOUR PARTICIPANTS WALKED AROUND

Area	Size ( $H \times W$ )	📷	MOTP ↓	MOTA ↑	IDF1 ↑	MT ↑	ML ↓	FPR ↓	FNR ↓	Rec1 ↑	Precn ↑	IDSR ↓	Frag ↓
A Activity Area	21m × 5m	7	1.38m	0.886	0.584	4	0	0.17	0.25	0.936	0.956	9.61	6
B Kitchen	3m × 24m	3	1.48m	0.671	0.427	1	0	0.28	1.01	0.747	0.913	17.38	11
C Left Corridor	33m × 2m	3	1.92m	0.545	0.435	1	0	0.27	1.53	0.616	0.901	12.97	4
D Lounge	3m × 8m	1	0.80m	0.334	0.501	0	2	0	2.66	0.334	1.0	0.0	0
E Right Corridor	33m × 1.5m	2	1.75m	0.602	0.468	2	0	0.01	1.56	0.610	0.993	13.10	9
F Staff Zone	14m × 3m	4	1.11m	0.734	0.659	3	0	0.65	0.40	0.899	0.845	4.44	1
Overall	-	-	<b>1.41m</b>	<b>0.629</b>	<b>0.512</b>	1.833	0.333	0.23	1.23	0.690	0.934	9.58	5.16
Five people	13m × 6m	4	1.44m	0.798	0.700	5	0	0.65	0.34	0.932	0.877	6.43	1

The area sizes for each section are shown with height and width in meters. 📷 Shows the number of cameras that were covering the movement over 85% of trajectories. Each area is annotated referring to figure 6. Overall performance shows average metrics across all areas combined. The last row additionally shows the quantitative results for simulating complex movements from five people with occlusions in activity area.

The bold number highlights the overall performance of the proposed system.

TABLE II  
BODY ORIENTATION ESTIMATION PERFORMANCE

Area	Size ( $H \times W$ )	📷	MAE- $\theta^\circ$	Acc.-5°	Acc.-15°	Acc.-22.5°	Acc.-30°	Acc.-45°	Acc.-90°
A Activity Area	21m × 5m	7	25.91°	0.319	0.622	0.723	0.800	0.855	0.912
B Kitchen	3m × 24m	3	31.06°	0.207	0.489	0.590	0.650	0.748	0.916
C Left Corridor	33m × 2m	3	21.68°	0.223	0.578	0.706	0.793	0.887	0.955
D Lounge	3m × 8m	1	47.37°	0.038	0.137	0.302	0.423	0.554	0.846
E Right Corridor	33m × 1.5m	2	32.41°	0.214	0.479	0.556	0.602	0.674	0.946
F Staff Zone	14m × 3m	4	33.68°	0.096	0.324	0.464	0.589	0.778	0.942
Overall	-	-	<b>28.97°</b>	0.220	0.506	0.615	0.693	0.789	0.930
Five people	13m × 6m	4	48.33°	0.149	0.328	0.417	0.512	0.629	0.814
In-place (Sess. 1)	11m × 5m	3	73.73°	0.087	0.181	0.255	0.261	0.322	0.644
In-place (Sess. 2)	4m × 7m	2	65.28°	0.078	0.212	0.236	0.290	0.375	0.672
In-place (Sess. 3)	3m × 5m	4	50.47°	0.111	0.265	0.454	0.580	0.692	0.769
In-place (Overall)	-	-	63.40°	0.091	0.218	0.310	0.371	0.457	0.693

The first six rows show the body orientation estimation performance at each region while four people walk in the space, and the seventh row shows the overall performance considering the entire area. 5 people moving with a more complex trajectory in activity area are shown in the eighth row. The last four rows show the estimation performance when a person stands in space but rotates only their orientations and the first three of these four rows show the evaluation from three different sessions. The last row shows the aggregated results.

The bold number highlights the overall performance of the proposed system.

## B. Body Orientation Estimation

The results of body orientation estimation are provided in Table II. Our evaluations encompass scenarios involving four individuals walking across the entire space and five individuals passing by each other in the activity area. In addition, we conducted assessments of body orientation for a scenario in

which five individuals remained stationary while altering their orientations within different group formations.

1) *Four People Walking Around Entire Space*: During the walking scenario with four participants, the overall MAE in body orientation estimation was 28.97°, and the accuracy within a 45° margin reached 78.9%, as indicated in the seventh row of Table II. Further examination of the correlation between



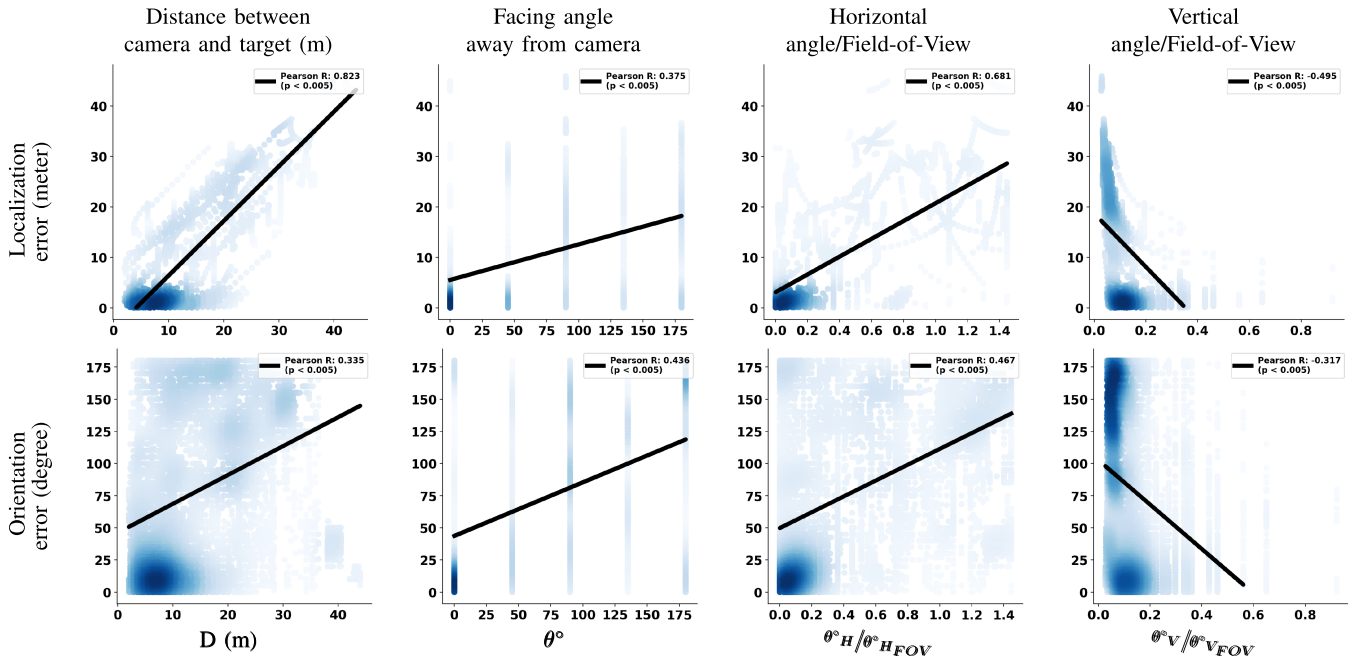


Fig. 10. Localization and orientation errors regarding camera installation. Each of the aspects is associated with localization (first row) and orientation estimation (second row) errors in each sample. The correlation plot shows denser samples with blue.

camera installation and body orientation estimation is presented in Section III-C.

2) *Five People Walking Past Each Other*: In situations where individuals were walking closely past each other, our body-orientation estimation technique faced challenges due to the presence of multiple people with different orientations captured within the same human bounding box. As illustrated in the eighth row of Table II, this posed difficulties, resulting in a high MAE of 48.33° and a lower accuracy within a 45° margin (62.9%).

3) *Five People Changing Orientations in Place*: The final four rows in Table II present the estimation performance when an individual remains stationary but rotates their orientation. In this scenario, the overall performance indicated an MAE of 63.40°.

### C. Impact of Parameters of Camera Installation

In Fig. 10, a Pearson correlation analysis is presented, highlighting the relationship between each factor from Fig. 8 and localization and orientation errors. In the figure, areas with denser samples are represented in a deeper shade of blue.

1) *Distance Between Camera and Person*: In Fig. 10, in the first column, a strong positive correlation is evident between the localization error (first row) and the camera distance, especially when the distance exceeds 20 m ( $R = 0.823$ ,  $p < 0.005$ ). Similarly, the orientation error (first column, second row) also exhibits a positive correlation with an increase in the distance between the captured individual and the camera ( $R = 0.335$ ,  $p < 0.005$ ).

2) *Facing Angle With Respect to Camera Viewpoint*: In Fig. 10, the second column, a strong positive correlation is

observed between the facing angle concerning the camera viewpoint and both localization ( $R = 0.375$ ,  $p < 0.005$ ) and body orientation estimation ( $R = 0.436$ ,  $p < 0.005$ ) errors.

3) *Horizontal Displacement From Camera Center*: In Fig. 10, in the third column, it is evident that the localization and body orientation estimation errors displayed significant increases, with strong positive correlations, as the captured individual moved away from the horizontal center. The correlation coefficients for localization and body orientation estimation errors were  $R = 0.681$  and  $R = 0.467$ , respectively, both with  $p < 0.005$ .

4) *Vertical Displacement From Camera Center*: As the individual moved closer to the vertical camera center, there was a notable reduction in both the localization and orientation estimation errors. This decrease in errors can be attributed to the increased coverage of more body parts within the camera's field-of-view. This relationship was marked by strong negative correlations, with correlation coefficients of  $R = -0.495$  for localization and  $R = -0.317$  for orientation estimation, both with  $p < 0.005$ .

## IV. DISCUSSION

Our experiment results demonstrated the possibility of using edge computing devices and cameras to continuously monitor fine-grained activities across large indoor spaces.

### A. Multiperson Tracking and Localization

1) *Four People Walking Around Entire Space*: The evaluations show that the proposed model consistently tracks individuals throughout the sequence as long as they are initially detected. Furthermore, the system exhibits an ability to handle a certain degree of missing detections (false negatives) arising from the

limitations of the multiperson pose detection model on edge computing systems. To run the pose detection model in real time on an edge computing system, which is a resource-constrained device, we are currently using a model that is specifically tailored to emphasize efficiency over performance compared to the state-of-the-art models [10]. This emphasis on efficiency is reflected in the system's high Prcn (93.48%), high MT (1.83), and low ML (0.33). The overall performance of localization remains within a reasonable error range, given our tracking rate of 1 Hz, which is adequate considering that people typically walk at a rate of 1.42 m/s in everyday activities [30], [31], [32]. In addition, the Prcn of our ground-truth annotations includes a margin of 1-m error since participants' locations were recorded as the nearest 1-m interval marker in the space.

The model's performance exhibited variations based on the regions within the study space. In the activity area (A), where most physical activities occur, the MOTA score was 88.63%, with a MOTP of 1.38 m, high Prcn (95.6%), and Rcll (93.62%), which are sufficient for detecting, tracking, and understanding group activities and space usage behavior in this area. In the activity area (A), seven cameras were tracking movements, emphasizing the importance of multiview cameras in noisy detection scenarios. The overlapping viewpoints of these cameras facilitated capturing people from various angles, including the front, back, and sides during walking, as demonstrated in Fig. 2.

On the other hand, lower MOTA scores were observed in regions with only a few cameras covering a wide capture space. Specifically, the kitchen, left corridor, and right corridor recorded MOTA scores of 67.16%, 54.53%, and 60.25%, with corresponding MOTP values of 1.48, 1.92, and 1.75 m, respectively. These regions featured long corridors with dimensions of 24 m  $\times$  3 m, 33 m  $\times$  2 m, and 33 m  $\times$  1.5 m, respectively. The limited camera coverage for these regions, and capturing individuals from a distance, led to challenges related to occlusion and pose detection. The small pixel region allocated for each person resulted in elevated FNRs of 1.01, 1.53, and 1.56, respectively. Nonetheless, our system exhibited consistently high Prcn (95.6%, 91.3%, 100%) in these regions, indicating that the detected objects were predominantly accurately localized.

The staff zone is equipped with four cameras positioned to capture movements from the sides (left and right) while walking. The size of the staff zone, which measures 14 m  $\times$  3 m, is the smallest among all areas, leading to cameras capturing individuals at closer distances. As a result, this area achieved a higher MOTA score of 73.47% and had a lower MOTP score of 1.11 m and reduced FNRs of 0.40 compared with other areas. Conversely, the lounge area only features a single camera covering the entire space and exhibited the lowest performance, with a MOTA score of 33.44%. This was primarily attributed to the lack of power and network resources near the lounge area, limiting the use of a single camera for coverage. In addition, occlusions from the front desk, chairs, and other objects in the lounge further impacted the tracking accuracy. In summary, our experimental results highlight the importance of multiview cameras positioned at closer distances to individuals. These factors are essential for addressing challenges associated with noisy pose estimations in edge computing devices, ultimately leading to lower MOTP scores and improved MOTA performance.

2) *Five People Walking Past Each Other*: Our findings indicate that employing four multiview cameras to capture activities

from different perspectives can help mitigate occlusions when individuals are in motion, even in the presence of noisy or missing pose estimations.

## B. Body Orientation Estimation

1) *Four People Walking Around Entire Space*: The body-orientation estimation performance reliably distinguished between eight orientations (N, NE, E, SE, S, SW, W, and NW). However, the performance of orientation estimation varied depending on the size of the captured area and the overlap between multiview cameras. In the activity area, which had seven cameras capturing views from different angles, the orientation estimation errors were relatively low, with an MAE of 25.91° (see Table II, first row). The staff zone, equipped with four cameras providing overlapping views and close proximity to individuals, also exhibited lower MAE (32.41°) (see Table II, sixth row).

Surprisingly, the left and right corridors, despite having low MOTA scores, demonstrated low-orientation estimation errors with MAEs of 21.68° and 32.41°, respectively (see Table II, third and fifth rows). These areas faced challenges in localization and tracking due to fewer cameras covering long and narrow distances. When estimating body orientations, the narrow width of the corridors made people tend to face forward or opposite to the camera viewpoints when they move, making it easier for the model to estimate body orientations. In contrast, the lounge area, covered by a single camera and experiencing occlusions from objects, found it challenging to identify body orientations. This resulted in a higher MAE of 47.37° and lower accuracy-45° of 55.4% (see Table II, fourth row). Overall, the placement of cameras is dependent on the space's structural characteristics. Fewer cameras (1 or 2) suffice for narrow corridors to capture linear movements, while open spaces benefit from multiview cameras to account for complex and unrestricted wandering.

2) *Five People Walking Past Each Other*: The body orientation estimation model struggled with occlusions when people were close to each other. However, it still managed to achieve an accuracy-90° score of 81.4%. This indicates that our model can reliably identify the four cardinal directions (N, E, S, and W) even when people are in close proximity to each other and changing group formations dynamically. These insights establish a foundation for analyzing group activities in crowded areas, as successful group activity analysis necessitates the identification of group formations and facing directions during movement.

3) *Five People Changing Orientations in Place*: In this scenario, our model was capable of identifying the four cardinal directions (N, E, S, and W) with a minimum accuracy of 69.3% (Acc.-90°). Similar to the previous experiment, the primary source of error was occlusions between people who were in close proximity, resulting in multiple individuals being captured in a single human bounding box. In addition, our body orientation estimation model struggled to detect subtle changes in facing direction (approximately 5°–10°) between consecutive frames, as it estimates orientations frame by frame. This indicates that our model performed better when individuals were walking, as our multiperson tracking method updated the orientation information based on their movement directions.

The accuracy of each session varied based on the number of cameras and the size of the areas they covered. In Session 1 (see

Table II, ninth row), which was covered by three cameras and had a space of 55 m<sup>2</sup>, the MAE was 73.7°. It is worth noting that the samples in this session consisted of smaller groups with only two or three members. Session 2 (see Table II, tenth row) had a slightly lower MAE of 65.2°. In this session, group sizes were larger (four members), but the covered area was smaller (28 m<sup>2</sup>) than in Session 1. This demonstrates the significance of multiview cameras for improved accuracy. Session 3 (see Table II, 11th row) was particularly noteworthy, as it had the lowest MAE (50.4°) among all three sessions. Despite having a larger social group of five people, the capture space was only 15 m<sup>2</sup>, and there were four cameras available, highlighting the importance of smaller areas and more cameras for better performance.

### C. Impact of Parameters of Camera Installation

1) *Distance Between Camera and Person*: The majority of our samples exhibited localization errors within approximately 2 m and orientation errors within 25° when the distance between the person and the camera was less than 15 m. This observation highlights that for the edge computing cameras, the portion of pixels corresponding to a person significantly diminishes as the person moves beyond 15 m. The models faced more challenges when individuals obscured each other, making it appear as if multiple people were a single person when viewed from a distance. In addition, the pose estimation models often struggled to distinguish between frontal and rear views of a person, particularly when the person was positioned over 20 m away from the camera, resulting in distorted keypoint predictions.

2) *Facing Angle With Respect to Camera Viewpoint*: The orientation at which the captured person was facing had a pronounced influence on the accuracy of both localization and body orientation estimation. Facing angles within 50° of the camera's line-of-sight was identified as pivotal in achieving a localization error of less than 2 m and a body orientation error of less than 25°. The 2-D pose estimation model exhibited its optimal performance when the person was facing the camera directly, resulting in more confident estimations.

3) *Horizontal Displacement From Camera Center*: To achieve a localization error of less than 2 m and an orientation error of less than 25°, it is essential that a person is within 30% of the horizontal field-of-view. When a person is captured too close to the frame boundaries, the lens distortion in the edge computing camera disrupts the planar assumption of the floor, leading to a degradation in the perspective transformation (as discussed in Section II-C3). This disruption significantly increases the localization error. In addition, when subjects are only partially within the frame boundaries, the orientation estimation error also escalates.

4) *Vertical Displacement From Camera Center*: As illustrated in Fig. 10 (fourth column), our edge computing camera module exhibited limitations when it came to capturing the foot locations of individuals with  $\theta_V^\circ / \theta_{V_{FOV}}^\circ \geq 0.8$  normalized vertical angle. In such cases, the pose estimation model failed to detect the person. On the contrary, when the value of  $\theta_V^\circ / \theta_{V_{FOV}}^\circ$  approached zero, both localization and body orientation errors showed a significant increase, primarily due to the greater distance between the captured person and the camera. The optimal

range for the normalized vertical angle was found to be between  $0.05 \leq \theta_V^\circ / \theta_{V_{FOV}}^\circ \leq 0.2$ . This corresponds to distances ranging from approximately 6 to 23 m between the captured person and the camera. This observation aligns with the conclusions drawn from the analysis of horizontal camera distance discussed earlier.

### D. Limitations

We identified structural variances of approximately 1.5 m compared with the building blueprint, which was used to develop our localization model in Section II-C3. Nevertheless, we deemed a margin of 2 m for localization error to be acceptable for effectively tracking participant behavior at our study site in future applications.

When it comes to annotating body orientation, achieving precise turns in exact intervals of 5°–10° proved to be challenging for our participants. However, we consider an error range of up to 45° to be reasonable for tracking social interactions during group activities. The normal horizontal field-of-view of the human eye is about 135°, and our system can detect dyadic interactions from overlapping field-of-view between two people with at most 90° margin of error. In our analysis, the MAE for body orientation was found to be 28.97° overall, which falls well within the acceptable 45° margin of error.

The areas that were not covered in our evaluation are filled with tables, yoga mats, or gym equipment. Future studies will examine how complex movements and occlusions affect our tracking and body orientation estimation methods when people interact with those furniture and equipment. Despite these limitations in our evaluation, the proposed work shows promise as an initial step toward analyzing complex movements in open spaces using a privacy-preserving camera-based activity monitoring system using secure edge computing framework.

### E. Future Applications

*Clinical Application*: Our system allows continuous and objective assessment of gait, spatial navigation, and group activities. Our system can be used to understand behavior and potentially detect behavioral changes in individuals with various motor and cognitive impairments [35]. The future work will further enable a system that may be able to predict falling episodes in real time, thereby preventing an adverse event, by directly using 2-D poses from multiview cameras across the facility [36], [37].

*Workspace Design*: Social interactions at work, which are highly correlated with productivity and well-being, are influenced by workplace design [38]. Our system can provide detailed information on socio-spatial formation patterns, such as space occupancy, behavioral mapping [39], Proxemics (physical distance between the people) [40], [41], and F-formation (shape and size of group formation) [42], [43], which allows the designer to take into account to design the workplace.

## V. CONCLUSION

Camera-based movement monitoring systems are gaining popularity thanks to the progress in edge computing platforms and computer vision technologies. In this work, we demonstrated the practicality of employing edge computing cameras



to localize, estimate body orientations, and track multiple individuals within a large indoor space while not compromising on privacy. Our sensing hub leveraged cost-effective edge computing components, including a Raspberry Pi, its camera module, a TPU, and a computing infrastructure that can be easily installed in pre-existing building setups.

In our study, the described system demonstrated highly reliable multiperson tracking (62.0%–88.6% MOTA) and body orientation estimation (29.0°) across various functional spaces. For the region with more cameras ( $7 >$ ), our system could track people in crowded activities with a higher level of accuracy (88.6% MOTA). In addition, we anticipate that our comprehensive examination of camera installation will offer valuable insights for individuals interested in implementing our system to enhance their environments. In our future work, we plan to increase the camera density based on the insights from our analysis. We will also assess the viability of tracking social interactions in more intricate settings, such as those involving group activities amid furnishings or the use of gym equipment.

## REFERENCES

- [1] A. Haque et al., “Towards vision-based smart hospitals: A system for tracking and monitoring hand hygiene compliance,” in *Proc. Mach. Learn. Healthcare Conf.*, 2017, pp. 75–87.
- [2] A. C. Yang, N. Lau, and J. C. Ho, “The role of bedroom privacy in social interaction among elderly residents in nursing homes: An exploratory case study of Hong Kong,” *Sensors*, vol. 20, no. 15, 2020, Art. no. 4101.
- [3] P. E. Lopez-de Teruel, F. J. Garcia, O. Canovas, R. Gonzalez, and J. A. Carrasco, “Human behavior monitoring using a passive indoor positioning system: A case study in a SME,” *Procedia Comput. Sci.*, vol. 110, pp. 182–189, 2017.
- [4] O. Dogan, J. Bayo-Monton, C. Fernandez-Llatas, and B. Oztaysi, “Analyzing of gender behaviors from paths using process mining: A shopping mall application,” *Sensors*, vol. 19, no. 3, 2019, Art. no. 557.
- [5] N. K. Vuong, S. Chan, and C. T. Lau, “Automated detection of wandering patterns in people with dementia,” *Gerontechnol.*, vol. 12, no. 3, pp. 127–147, 2014.
- [6] I. Cheol Jeong et al., “Using a real-time location system for assessment of patient ambulation in a hospital setting,” *Arch. Phys. Med. Rehabil.*, vol. 98, no. 7, pp. 1366–1373, 2017.
- [7] T. Van Haute et al., “Performance analysis of multiple indoor positioning systems in a healthcare environment,” *Int. J. Health Geographics*, vol. 15, no. 1, pp. 1–15, 2016.
- [8] S. Yoo, S. Kim, E. Kim, E. Jung, K. Lee, and H. Hwang, “Real-time location system-based asset tracking in the healthcare field: Lessons learned from a feasibility study,” *BMC Med. Informat. Decis. Mak.*, vol. 18, no. 1, pp. 1–10, 2018.
- [9] Y. Dai, Z. Hu, S. Zhang, and L. Liu, “A survey of detection-based video multi-object tracking,” *Displays*, vol. 75, 2022, Art. no. 102317.
- [10] E. S. dos Reis et al., “Monocular multi-person pose estimation: A survey,” *Pattern Recognit.*, vol. 118, 2021, Art. no. 108046.
- [11] T. Tsai, C. Chang, and S. Chen, “Vision based indoor positioning for intelligent buildings,” in *Proc. 2nd Int. Conf. Intell. Green Building Smart Grid*, 2016, pp. 1–4.
- [12] A. Cosma, I. E. Radoi, and V. Radu, “CamLoc: Pedestrian location estimation through body pose estimation on smart cameras,” in *Proc. Int. Conf. Indoor Positioning Indoor Navigation*, 2019, pp. 1–8.
- [13] Z. Luo et al., “Computer vision-based descriptive analytics of seniors’ daily activities for long-term health monitoring,” *Mach. Learn. Healthcare*, vol. 2, no. 1, 2018.
- [14] D. Xue et al., “Vision-based gait analysis for senior care,” 2018, *arXiv:1812.00169*.
- [15] X. Jin, L. Li, F. Dang, X. Chen, and Y. Liu, “A survey on edge computing for wearable technology,” *Digit. Signal Process.*, vol. 125, 2022, Art. no. 103146.
- [16] P. B. Suresha, C. Hegde, Z. Jiang, and G. D. Clifford, “An edge computing and ambient data capture system for clinical and home environments,” *Sensors*, vol. 22, no. 7, 2022, Art. no. 2511.
- [17] J. Barthélemy, N. Verstaevl, H. Forehead, and P. Perez, “Edge-computing video analytics for real-time traffic monitoring in a smart city,” *Sensors*, vol. 19, no. 9, 2019, Art. no. 2048.
- [18] G. Papandreou, T. Zhu, L. Chen, S. Gidaris, J. Tompson, and K. Murphy, “PersonLab: Person pose estimation and instance segmentation with a bottom-up, part-based, geometric embedding model,” in *Proc. Eur. Conf. Comput. Vis.*, 2018, pp. 269–286.
- [19] T. Lin et al., “Microsoft COCO: Common objects in context,” in *Proc. 13th Eur. Conf. Comput. Vis.*, 2014, pp. 740–755.
- [20] X. Li, K. Wang, W. Wang, and Y. Li, “A multiple object tracking method using Kalman filter,” in *Proc. IEEE Int. Conf. Inf. Autom.*, 2010, pp. 1862–1866.
- [21] R. E. Kalman, “A new approach to linear filtering and prediction problems,” *J. Basic Eng.*, vol. 82, no. 1, pp. 35–45, Mar. 1960.
- [22] W. Chan and F. Hsiao, “Implementation of the Rauch-Tung-Striebel smoother for sensor compatibility correction of a fixed-wing unmanned air vehicle,” *Sensors*, vol. 11, no. 4, pp. 3738–3764, 2011.
- [23] M. Sun, L. Zhang, Y. Liu, X. Miao, and X. Ding, “See-your-room: Indoor localization with camera vision,” in *Proc. ACM Turing Celebration Conf.-China*, 2019, pp. 1–5.
- [24] C. Wu et al., “MEBOW: Monocular estimation of body orientation in the wild,” in *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit.*, 2020, pp. 3451–3461.
- [25] Z. Zhang, “A flexible new technique for camera calibration,” *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 22, no. 11, pp. 1330–1334, Nov. 2000.
- [26] C. John and H. D. Allan, *A First Look at Graph Theory*. Singapore: Allied Publishers, 1995.
- [27] H. W. Kuhn, “The Hungarian method for the assignment problem,” *Nav. Res. Logistics Quart.*, vol. 2, no. 1/2, pp. 83–97, 1955.
- [28] K. Bernardin, A. Elbs, and R. Stiefelwagen, “Multiple object tracking performance metrics and evaluation in a smart room environment,” in *Proc. 6th IEEE Int. Workshop Vis. Surveill., conjunction ECCV*, Citeseer, vol. 90, no. 91, May 2006.
- [29] P. Dendorfer et al., “MOT20: A benchmark for multi object tracking in crowded scenes,” 2020, *arXiv:2003.09003*.
- [30] R. C. Browning, E. A. Baker, J. A. Herron, and R. Kram, “Effects of obesity and sex on the energetic cost and preferred speed of walking,” *J. Appl. Physiol.*, vol. 100, no. 2, pp. 390–398, 2006.
- [31] B. J. Mohler, W. B. Thompson, S. H. Creem-Regehr, H. L. Pick, and W. H. Warren, “Visual flow influences gait transition speed and preferred walking speed,” *Exp. Brain Res.*, vol. 181, no. 2, pp. 221–228, 2007.
- [32] R. V. Levine and A. Norenzayan, “The pace of life in 31 countries,” in *J. Cross-cultural Psychol.*, vol. 30, no. 2, pp. 178–205, 1999.
- [33] D. Yu, H. Xiong, Q. Xu, J. Wang, and K. Li, “Continuous pedestrian orientation estimation using human keypoints,” in *Proc. IEEE Int. Symp. Circuits Syst.*, 2019, pp. 1–5.
- [34] M. Andriluka, S. Roth, and B. Schiele, “Monocular 3D pose estimation and tracking by detection,” in *Proc. IEEE Comput. Soc. Conf. Comput. Vis. Pattern Recognit.*, 2010, pp. 623–630.
- [35] M. Lindh-Rengifo et al., “Components of gait in people with and without mild cognitive impairment,” *Gait Posture*, vol. 93, pp. 83–89, 2022.
- [36] H. W. Lach, B. E. Harrison, and S. Phongphanngam, “Falls and fall prevention in older adults with early-stage dementia: An integrative review,” *Res. Gerontological Nurs.*, vol. 10, no. 3, pp. 139–148, 2017.
- [37] G. Allali and J. Verghese, “Management of gait changes and fall risk in MCI and dementia,” *Curr. Treat. Options Neurol.*, vol. 19, pp. 1–18, 2017.
- [38] B. Lee, M. Lee, P. Zhang, A. Tessier, and A. Khan, “An empirical study of how socio-spatial formations are influenced by interior elements and displays in an office context,” in *Proc. ACM Hum.-Comput. Interact.*, vol. 3, pp. 1–26, 2019.
- [39] A. Cox, J. Loebach, and S. Little, “Understanding the nature play milieu: Using behavior mapping to investigate children’s activities in outdoor play spaces,” *Child. Youth Environ.*, vol. 28, no. 2, pp. 232–261, 2018.
- [40] T. Ballendat, N. Marquardt, and S. Greenberg, “Proxemic interaction: Designing for a proximity and orientation-aware environment,” in *Proc. ACM Int. Conf. Interactive Tabletops Surfaces*, 2010, pp. 121–130.
- [41] E. T. Hall, “A system for the notation of proxemic behavior,” *Amer. Anthropologist*, vol. 65, no. 5, pp. 1003–1026, 1963.
- [42] M. Danninger, R. Vertegaal, D. P. Siewiorek, and A. Mamuji, “Using social geometry to manage interruptions and co-worker attention in office environments,” in *Proc. Graph. Interface*, 2005, pp. 211–218.
- [43] A. Kendon, “Spatial organization in social encounters: The F-formation system,” *Conducting Interaction: Patterns Behav. Focused Encounters*, 1990.