LiDArgus: Accommodating Diverse Privacy Preferences in LiDAR Sensing Through Configurable Sub-Zoning

BoRui Li
School of Computing Science
Simon Fraser University
Burnaby, British Columbia, Canada
bla127@sfu.ca

Jinda Yang
School of Computing Science
Simon Fraser University
Burnaby, British Columbia, Canada
jinda_yang@sfu.ca

Apurva Jain
Lumotive Technologies Inc.
Burnaby, British Columbia, Canada
apurva.jain@lumotive.com

William Newberry
Lumotive Technologies Inc.
Burnaby, British Columbia, Canada
william.newberry@lumotive.com

Xing-Dong Yang
School of Computing Science
Simon Fraser University
Burnaby, British Columbia, Canada
xingdong_yang@sfu.ca

Abstract

LiDAR sensors are increasingly being adopted in smart home environments to enable user activity recognition for a wide range of applications, including resident safety monitoring, health tracking, and behavior modeling. While the advantages of LiDAR-based activity recognition are well recognized, supporting diverse privacy preferences remains an ongoing challenge. Different user groups often have competing needs: homeowners may favor high-resolution sensing for health monitoring, whereas guests may prefer lower-resolution sensing or even temporary deactivation of the system. To navigate these varied requirements, we introduce a spatially adaptive approach that divides the LiDAR sensing field into configurable sub-zones, each with adjustable sensing capabilities, sizes, and locations. This design empowers users to personalize sensing settings according to their individual privacy preferences. We demonstrate the feasibility of this approach through a proof-of-concept prototype and assess its effectiveness via two user studies. Our findings offer insights to inform the design of privacy-aware LiDAR sensing systems in smart environments.

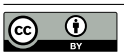
CCS Concepts

• **Human-centered computing** → **Interactive systems and tools**.

Keywords

LIDAR, Privacy, Activity Sensing

ACM Reference Format:
BoRui Li, Jinda Yang, Apurva Jain, William Newberry, and Xing-Dong Yang. 2025. LiDArgus: Accommodating Diverse Privacy Preferences in LiDAR Sensing Through Configurable Sub-Zoning. In *The 38th Annual ACM Symposium on User Interface Software and Technology (UIST '25), September 28–October 01, 2025, Busan, Republic of Korea.* ACM, New York, NY, USA, 23 pages. https://doi.org/10.1145/3746059.3747724



This work is licensed under a Creative Commons Attribution 4.0 International License.
UIST '25, Busan, Republic of Korea
© 2025 Copyright held by the owner/author(s).
ACM ISBN 979-8-4007-2037-6/25/09
https://doi.org/10.1145/3746059.3747724

1 INTRODUCTION

LiDAR sensors are increasingly being integrated into smart home environments to support user input detection and activity monitoring. They enable a wide range of applications, such as controller-free computer interactions [25, 57], resident safety monitoring [23, 44, 45, 64], and the collection of behavior and health data for data-driven services [22, 24, 40, 56, 68].

While the advantages of sensing are well recognized, the home remains a deeply personal space where privacy preferences can vary widely. These preferences may differ not only across individuals with different roles, such as space owners and visitors, but also depending on situational context. For instance, a homeowner might appreciate the detection of health-related behaviors, such as drinking water, while a visitor may prefer not to have their actions sensed to avoid potential observation by others later [42, 72]. Supporting such diverse preferences in a shared space, without introducing conflict, presents an ongoing challenge [42, 72]. Even among the same group of users, privacy preferences can shift based on location within the home [32]. In more relaxed settings, like a living room couch, users may prefer greater privacy, particularly during socially awkward moments (e.g., scratching their back), to avoid the possibility of these actions being recorded or later viewed by others [10, 42]. To ensure complete privacy, they may opt to disable the sensor. However, in nearby work areas like a desk, users may expect productive activities to be sensed and supported. This requires frequent sensor activation and deactivation as the user moves between locations, which can be cumbersome.

In this paper, we introduce a method designed to better accommodate diverse privacy preferences in LiDAR-based activity recognition applications. Our approach divides the LiDAR's sensing area into configurable sub-sensing zones with adjustable resolutions. We show that even a simple partitioning, such as along the x-axis in our implementation (Figure 1a), can effectively support a broader range of user privacy needs. Each sub-sensing zone can be independently toggled on or off and customized in terms of size, location, and sensing resolution to reflect individual privacy preferences. For example, a visitor might remain within a low-resolution zone that dynamically follows their movements, preventing their activity data from being identifiable to either humans or machine learning

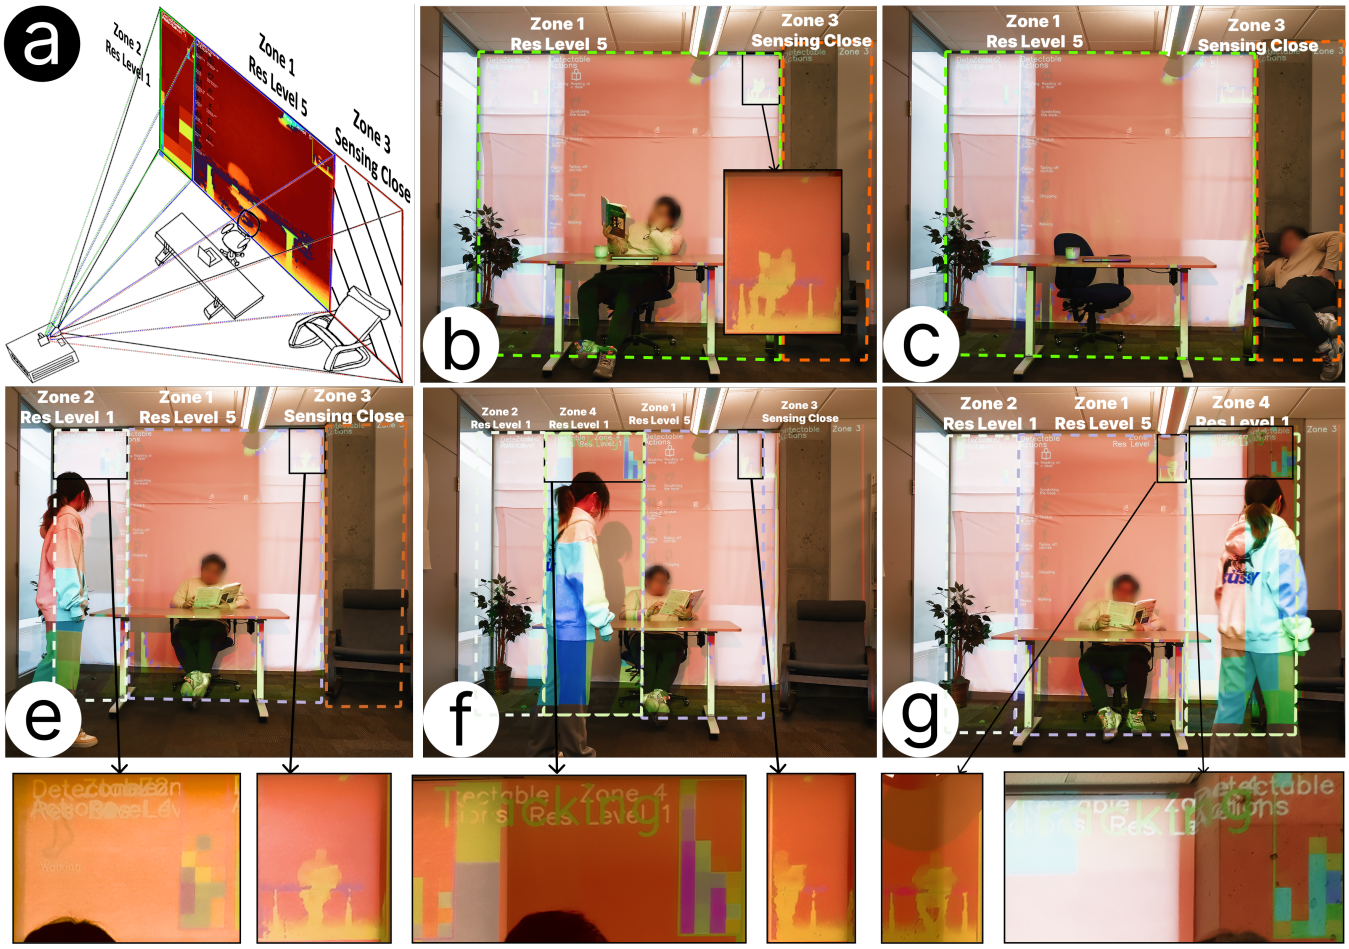


Figure 1: (a) LiDArgus improves users' privacy preferences in LiDAR activity recognition systems by dividing the LiDAR sensing area into configurable sub-sensing zones (three shown in this figure). Each zone has adjustable sensing capabilities, sizes, and locations, allowing users to tailor the sensing settings to their privacy needs. (b) Bob has set up a LiDAR activity recognition system in his new home to track his daily activities, supporting his professional management. His desk activities are monitored within a high-resolution zone, allowing the system to accurately recognize tasks such as sitting and reading. (c) To preserve privacy, Bob's rock chair is placed in a separate zone where sensing is disabled, ensuring that socially sensitive actions, such as scratching his back or picking his nose, are not captured. (d) When Bob's friend Alice visits, she creates a sub-sensing zone with the lowest resolution, acting as a privacy mask that follows her movements. (e-f) Alice's low-resolution zone ensures her activities remain unidentifiable to Bob's system's machine learning models or other people. While Alice perceives the room as being sensed at a lower resolution, areas outside her zone are maintained at a higher resolution to support Bob's activity recognition needs.

models (Figure 1d-f). At the same time, other areas of the space can retain high-resolution sensing to support the activity recognition needs of the space owner.

This flexible configuration allows the system to simultaneously uphold both privacy and sensing requirements for users in different roles. In scenarios where privacy preferences vary across different locations of the environment, a sub-sensing zone can also be used to fully disable sensing in specific areas, such as over a couch, creating a privacy-preserving region distinct from the rest of the space (Figure 1b-c). Overall, our method simulates the effect of deploying multiple LiDAR sensors with different resolutions and blind spots, while offering a more scalable and efficient solution using a single LiDAR sensor.

To validate our approach, we developed a proof-of-concept prototype, LiDArgus, using the M30 LiDAR Development Kit from Lumotive [2]. This device incorporates a solid-state beam steering mechanism powered by a Light Control Metasurface (LCM™), enabling sub-zoning directly at the hardware level. This capability allows selective control of the LiDAR's beam illumination and sampling, eliminating the need to first capture the full dataset and discard unnecessary information.

Our prototype supports the creation of dynamic sub-sensing zones along the x-axis to accommodate varying privacy needs. To enhance system transparency and user understanding, we integrated a projector that visualizes key information, such as zone locations, boundaries, sensing resolution, and the types of activities recognizable by the system's machine learning model. The zone boundaries are displayed within the physical space, aligning with their actual locations, and raw sensor data is projected onto the user's body and surrounding area. Users can create, adjust, and personalize sensing zones using voice commands, allowing them, including owners and visitors, to configure their zones freely, ensuring their privacy preferences are respected without interfering with those of others, or with their own preferences across different areas of the sensed space.

We conducted two user studies to assess the usability and effectiveness of LiDArgus. The first study investigated how effectively LiDArgus's visual feedback communicated information about sub-sensing zones during user interactions. Participants found the visual cues to be intuitive and easy to understand. The second study examined whether the introduction of sub-sensing zones, along with granting users control over these zones, could better support users in managing their privacy preferences across different roles. Our findings demonstrate that LiDArgus increased participants' awareness of the LiDAR sensor's operational state and improved perceived transparency regarding its coverage and sensing capabilities. Furthermore, by offering meaningful control over the sensor, LiDArgus more effectively supported users' privacy preferences compared to a baseline condition without sub-zoning.

2 CONTRIBUTION

The primary contribution of this work is the introduction of sub-zoning within a LiDAR sensing region as an approach to address the diverse privacy preferences of users in activity recognition applications. Rather than focusing on the technical implementation of enabling sub-zoning, our contribution emphasizes a new privacy-preserving user experience in LiDAR powered smart environments. Unlike existing methods that typically rely on a single sensing area with uniform resolution, our approach utilizes a novel LCM-based LiDAR system to create customizable sub-sensing zones to (1) provide transparency regarding the data collected by the LiDAR and the functionality of each sensing zone, and (2) empower users in different roles to create and configure the resolution, location, size, and behavior of their designated zones based on their specific privacy preferences.

3 RELATED WORK

We review related work on activity recognition using RGB cameras and LiDAR sensors, including efforts to enhance their privacy-preserving capabilities. Additionally, we examine prior research on understanding diverse privacy preferences in smart environments, which served as a key inspiration for our study.

3.1 Privacy-Preserving Activity Recognition in Camera-Based Systems

Activity recognition is a cornerstone of smart environments, enabling a wide range of applications, including home surveillance [48, 50, 53, 71], video monitoring [4, 8, 21, 39], and healthcare systems [6, 9, 15, 55]. Much prior work has focused on vision-based methods, using cameras to identify activities [27, 49, 61, 65]. However, the increasing use of these systems has raised important concerns regarding user privacy. Prior research has shown that the presence of cameras in smart environments can lead to discomfort and diminished trust, especially when users are unsure about what is being recorded and how the data is being utilized [67, 74]. In response to these concerns, researchers have proposed a range of privacy-preserving techniques. For instance, some approaches employ methods such as blurring sensitive areas in images [29, 46, 47] or applying mosaic filters to facial regions to support anonymization [19, 28, 34, 52, 60, 62]. Others aim to enhance activity recognition models to operate effectively with low-resolution images, thereby reducing the visibility of identifiable features while maintaining satisfactory recognition performance [7, 43, 59, 69, 70]. Despite these advances, a fundamental challenge remains: visual data is inherently rich in detail, and even with obfuscation or reduced resolution, contextual or biometric cues may still be inadvertently revealed. As such, ensuring complete privacy continues to be an ongoing area of exploration and innovation.

3.2 Activity Recognition Using LiDAR

Unlike cameras, LiDAR-based sensing does not capture appearance or texture details. Consequently, it avoids recording identifiable visual features such as skin tone, eye color, or iris patterns, making it a more privacy-preserving option for activity recognition [51, 54, 58, 63]. Prior work has explored a range of techniques to use LiDAR data for activity recognition. For instance, Glandon et al. proposed a method that recognizes user activities by extracting skeletal information from the sensor's point cloud data [20]. In a complementary approach, Zhou et al. employed a voxel-based representation, leveraging a 4D tensor instead of raw point clouds to improve model training efficiency[75]. In practical applications, Kodur et al. [33] employed LiDAR on a hospital robot to detect patient gait and pedestrian movement, facilitating collision avoidance. Joo et al. [30] deployed LiDAR in care facilities to monitor the activities of Alzheimer's patients, providing alerts for abnormal behaviors. Additionally, Fraccaro et al. [18] integrated LiDAR with environmental data to develop a remote monitoring system for elderly home care.

LiDAR and camera-based systems share a common limitation found in most IoT sensors: they typically use a uniform sensing resolution across the entire sensing region. Although overall resolution may be adjusted dynamically, it is generally not possible to independently configure resolution for specific subregions within the sensing field. This results in a "one-size-fits-all" sensing model, where the entire space is either sensed at a high resolution, which may compromise privacy, or at a low resolution, which could limit functionality. In this work, we propose that subdividing the LiDAR sensing region into smaller zones with independently configurable sensing capabilities offers a more flexible and effective solution for balancing privacy and functionality requirements.

3.3 Diverse Privacy Preferences in Smart Environments

A critical aspect often overlooked in the design of smart sensing systems, including those based on LiDAR, is the variability of privacy needs across different user groups and contexts. Research has demonstrated that individuals' privacy preferences and expectations are multifaceted, shaped by elements such as the purpose of data collection, the context in which data is collected, and the parties with whom the data is shared [73]. Privacy preferences also differ by role: for instance, visitors typically feel more comfortable with low-fidelity sensors [35], while owners of smart environments may have greater trust in device manufacturers and may not prioritize privacy protection, potentially resulting in inadvertent disclosures of privacy-sensitive information [72]. The role of context is particularly significant for owners, with factors such as location playing a key role in their adoption of privacy settings and shaping their perceptions of privacy [32].

Studies have shown that visitors are often unaware that sensor data collection may pose privacy concerns [42, 72]. Although they may not directly benefit from IoT devices, their data is still being gathered. One study found that visitors desire the ability to evaluate the current status of sensors and understand how their data is being captured. Additionally, they seek clarity on the purpose of data collection, its management, and any associated security risks [42]. Visitors often feel powerless when it comes to configuring sensors to align with their privacy preferences. Even when provided with control, they may hesitate to make adjustments due to social pressures or discomfort [16].

To mitigate these challenges, privacy assistants (PAs) have been proposed as a promising solution to help users discover and manage the data collection practices of nearby IoT devices [12–14]. For example, Fernandez et al.'s system [5] enables users to view the data being collected under current settings (e.g., facial features, eye movement). Research indicates that PAs can enhance user comfort in new smart sensing environments [16, 36]. However, existing PAs are not directly applicable to LiDAR systems, nor do they effectively address the "one-size-fits-all" challenge, as highlighted in the previous subsection.

Our work extends existing work by introducing the concept of sub-zoning to tackle the limitations of one-size-fits-all solutions. The design of our interactions is informed by prior studies on user privacy preferences, as discussed earlier in this section. Specifically, we cater to the diverse privacy preferences of LiDAR users by ensuring transparency about the data being collected, enabling informed decision-making, and providing users with control over the system to adjust the sensing capabilities of sub-zones based on their privacy preferences.

4 WALKTHROUGH

We demonstrate the new user experience enabled by our approach in a multifunctional room that serves as both a home office and a living room for hosting guests. This setup is common among young professionals, though the concept is also applicable to larger living spaces as well.

Background: Bob recently moved into this room and configured it to monitor his daily activities, supporting both his health and professional development (Figure 1). His desk activities are tracked within a high-resolution zone (like a 1D peephole within the full sensing range, with a resolution of 640×480), enabling the system to accurately detect tasks such as sitting and reading. Given his busy schedule, Bob often eats at the desk, and the system also detects eating and drinking activities. To maintain privacy, Bob's rock chair is located in a separate zone where sensing is disabled, ensuring that socially sensitive behaviors, such as scratching his back or picking his nose, are not captured. The zones can be visually represented in the physical space, with sensing pixels projected onto his body or the surrounding walls.

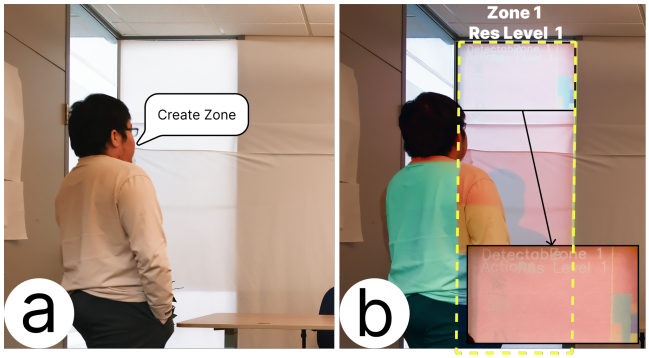


Figure 2: (a) Bob uses the voice command "Create Zone" to create a low-resolution private zone for his visitors. (b) LiDArgus generates a sub-sensing zone centered around Bob's location and visualizes information about the zone, including raw sensing data, sensing resolution, and detectable activities.

Scenario 1: Interacting with Zones. Bob's visitors often linger near the door for brief conversations. However, they feel uneasy knowing that their movements and activities are being detected by Bob's activity sensor. To address this concern, Bob decides to establish a private zone near the door. Standing at the entrance, he commands, "create zone," prompting the system to generate a designated area around him (Figure 2). This zone initially operates at the lowest sensing resolution (like a 1D peephole within the full sensing range, with a resolution of 40×30), allowing the machine learning model to detect only large movements, such as walking or falling, which are critical for safety, while excluding finer details to preserve privacy. To ensure that even human observers cannot easily interpret the raw data, Bob performs a few playful gestures and inspects the sensor output. The resolution appears sufficiently low that he cannot even recognize his own movements, leaving him confident in the system's ability to safeguard privacy. Although he could remove the zone at any time by saying, "Remove zone", Bob decides to keep it. A subtle sound effect provides audible confirmation whenever the zone is active.

This low-resolution zone near the door also serves as a designated "safe" area, providing Bob with a momentary retreat from the higher-resolution sensing zones that cover the rest of the room. After a busy Zoom meeting, Bob briefly steps into this zone to stretch his back, ensuring that this activity remains outside the scope of the sensors (see Figure 3). While Bob knows he can issue

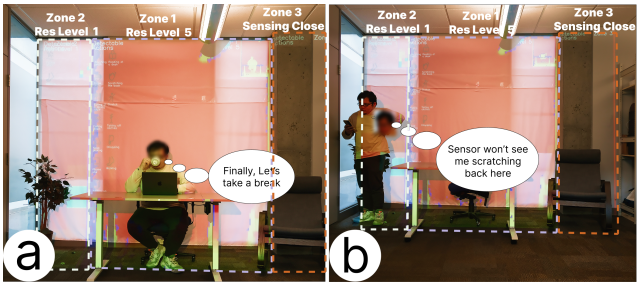


Figure 3: (a) Bob takes a break after a long Zoom meeting. He walks away from his desk and (b) steps into the new private zone, where he can perform activities, such as stretching his back, that he prefers not to be sensed by the LiDAR.

the voice command "minimize resolution" at his desk to reduce the workspace's resolution to its minimum, he prefers to step into the safe zone for convenience, as it is already within reach. Upon entering the zone, Bob expects to be informed of the new sensing area, a common preference for individuals transitioning into new sensor zones [13, 16, 42]. The system detects his movement and activates the projector to display both the zone boundaries and sensor data.

Scenario 2: Visitor. Alice, a friend of Bob, hears about his recent move and decides to visit him. As she enters, she unknowingly steps into the newly established private zone in Bob's room. The system does not differentiate between individuals but is able to detect the presence of a new person. This detection triggers a real-time visualization of the boundaries of all zones in the room (Figure 4a, b), offering immediate feedback on their spatial coverage. This design is guided by previous research highlighting the significance of user awareness of sensor presence [13].

The system also provides a list of detectable activities based on the resolution of each zone, enabling Alice to assess its capabilities and identify potential privacy concerns [42] (Figure 4b). Additionally, she can access the raw sensor data, which offers a clearer understanding of the system's operation in her current zone, further enhancing transparency, as recommended by prior research [16, 73]. After two minutes of inactivity, the zone visualization fades, but it can easily be reactivated with the voice command, "show all zones."

Curious about how raw sensor data appears in the high-resolution zone adjacent to the private zone, Alice extends her arm into this region. She quickly observes a much sharper and more detailed contour of her arm, offering an intuitive grasp of the resolution difference between the two zones (Figure 4c).

Alice decides to explore Bob's room. Before leaving the private zone, she creates her own by issuing the voice command "create zone" (Figure 5). Similar to Bob's guest zone, this new zone has a low sensing resolution, detecting only activities such as walking or falling. While a default resolution is helpful, Alice is interested in exploring the additional resolution options available in Bob's system. This desire for greater control is typical among individuals visiting new smart environments [42, 72]. She adjusts the resolution of her zone by using voice commands, such as "increase resolution," which increments it to the next predefined level. Alice also tries to disable sensing by saying "disable sensing," which turns off the sensors within her zone. After some experimentation, she determines that the default setting works best for her and opts to keep it. As shown in prior research [13], the ability to adjust sensor settings helps reassure her that her privacy is being protected (additional evidence supporting this is presented in our study results).

When Alice is ready to move, she activates the voice command "follow me," transitioning her zone from a static to a dynamic mode that continuously tracks her movements. This change enables the zone to adapt to her location in real time. As a result, Alice perceives the entire room as being sensed at a lower resolution, while areas outside her zone remain at a higher resolution to support Bob's needs for activity recognition (Figure 6).

In our scenario, Bob does not have a dynamic zone that tracks his movements. If such a zone were present, it could overlap with Alice's zone when they stand close together, resulting in the two zones merging into a larger one. In this case, the system would automatically adjust the sensing resolution to the lower of the two zones, prioritizing the more restrictive privacy settings. The projector would update the zone boundaries and resolution to notify the user (Figure 7). Once Alice leaves, her zone would be automatically removed.

5 IMPLEMENTATION

This section provides the implementation details of our prototype.

5.1 Hardware

The core component of our hardware is the M30 LiDAR Development Kit from Lumotive [2] (Figure 8). This device incorporates solid-state beam steering based on Light Control Metasurface (LCM™) technology, which enables sub-millisecond steering of the LiDAR beam without relying on mechanical components. It offers a horizontal and vertical field of view of $120° \times 90°$, respectively. LCM™ allows for precise, random-access control of scanning patterns, enabling the LiDAR sensor to selectively exclude specific areas within its sensing region by steering the laser beam to skip certain angles during environmental scanning. By adjusting the interval between adjacent beam scanning angles, the sensing resolution can be adjusted. Sensing can be effectively disabled in areas where beam scanning is skipped. This functionality enables real-time, no-latency sub-zoning at the hardware level. In addition to the LiDAR system, the hardware setup includes a Sony VPL-PHZ50 projector, used to display visual information about the sub-sensing zones directly onto the physical environment. The LiDAR's raw sensor data is projected onto the user's body. Audio feedback is delivered through Sony XE300 speakers.

5.2 Software

We used the Lumotive SDK to control the LiDAR's beam steering via its Light Control Metasurface and to acquire data from the device. To enable sub-zoning within the M30's sensing region, we used its virtual sensor functionality, which allows programmable control of beam steering to define multiple sub–field-of-view regions, each with different scanning resolutions and frame rates. Data communication was handled via TCP. To generate the raw sensing data presented to the user, we projected the point cloud

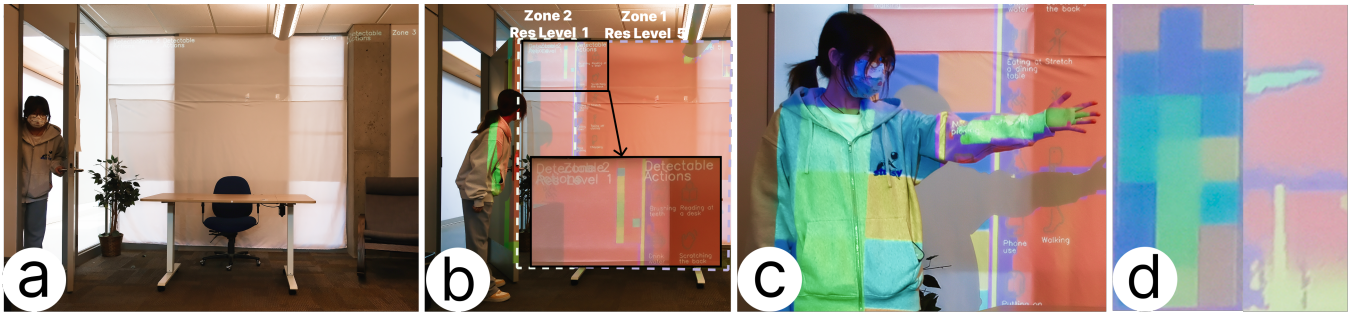


Figure 4: (a) Bob's friend Alice comes to visit. As she enters the room, she steps into the private zone. This detection triggers a visualization of all zone boundaries in the room, providing immediate feedback on their spatial coverage. (b) The system also presents a list of detectable activities based on each zone's sensing resolution, enabling Alice to assess the system's capabilities and identify potential privacy concerns. (c) She can see raw sensor data projected on her body to better understand how the system operates in her current zone. When Alice extends her arm into the high-resolution zone adjacent to the private zone, she sees a much sharper and more detailed contour of her arm, which conveys the resolution difference between the two zones (see the larger sensor pixels on her body vs smaller ones on her arm). (d) A thumbnail of the sensor data shows Alice's body being sensed at the lowest resolution and her arm at the highest resolution.

Resolution Level				
640×480 (Resolution Level 5)	320×240 (Resolution Level 4)	160×120 (Resolution Level 3)	80×60 (Resolution Level 2)	40×30 (Resolution Level 1)
Brushing teeth	Brushing teeth			
Drink water				
Eating at a dining table	Eating at a dining table			
Falling down	Falling down	Falling down	Falling down	Falling down
Reading at a desk	Reading at a desk			
Chopping				
Walking	Walking	Walking	Walking	Walking
Nose picking	Nose picking	Nose picking		
Phone use	Phone use	Phone use		
Putting on clothes				
Taking off clothes				
Stretch	Stretch	Stretch		
Scratching the back	Scratching the back	Scratching the back		

Table 1: Detectable activities under different resolutions in LiDArgus (based on yolo11m-cls with f1 score > 90%)

data onto a 2D plane. These projections were calibrated to ensure accurate alignment of raw sensor data with the user's body and precise mapping of zones to their physical locations in the environment. Voice input for detecting user commands was implemented using Google Speech-to-Text API [1].

5.3 Machine Learning

We trained a machine learning model to recognize a set of daily activities commonly encountered in home and office environments.

Activity Selection. We selected two groups of activities based on their varying levels of privacy sensitivity in smart home environments, as identified in prior research [10]. The first group consisted of activities generally considered to have low or no privacy concerns, including: (1) brushing teeth, (2) drinking water, (3) eating at a dining table, (4) falling down, (5) reading at a desk, (6) chopping, and (7) walking. The second group included activities that are often regarded as privacy-sensitive, including: (1) taking off clothes, (2) putting on clothes, (3) nose picking, (4) phone use, (5) stretching, and (6) scratching the back. The inclusion of privacy-sensitive activities was to help participants envision the types of behaviors that the system might detect. In practice, this data may not be collected at all.

Data Collection. Many of the selected activities are not included in existing open-source human activity depth map datasets[17, 37, 38]. Therefore, we collected our own dataset. To do so, we adopted the data collection methodology used in the NTU120 dataset [37]. During data collection, we positioned the M30 LiDAR Development Kit 2.5m above the ground and 3.5m away from the participants. For each activity, we recorded sensor data with a full FoV of $120° \times 90°$ at multiple sensing resolutions, ranging from the system's highest supported resolution of 680×480 down to 320×240, 160×120, 80×60, and 40×30. A total of 12 participants took part in the data collection, with heights ranging from 163 cm to 185 cm. For each activity performed by a participant, we recorded 60 consecutive frames. After all actions were completed, we also captured additional background data to serve as negative samples.

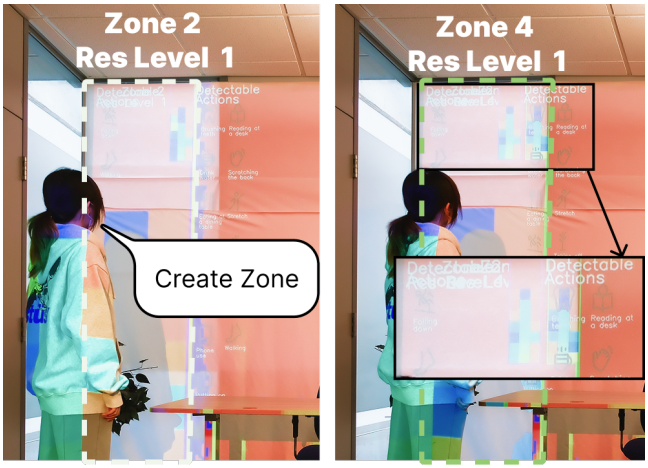


Figure 5: (a) Alice decides to explore Bob's room. Before leaving her private zone, she issues the voice command "create zone" to create a new one. (b) The newly created zone overlaps with Bob's private zone near the door and operates at the lowest sensing resolution (see the large sensor pixel on her body).

Model Selection. Given that much of the privacy-sensitive data is processed on edge devices, we focused our model selection on approaches capable of real-time inference under limited computational resources. Although our data exhibits temporal characteristics, existing sequence-based prediction models are not suitable for real-time inference. Therefore, we treated the activity recognition task as a classification problem and adopted YOLO11 for its efficiency. We used object detection to determine the body's position and size within the sensor data, which is crucial for adjusting the size and location of the sub-sensing zones.

Data Processing and Model Training. We used the YOLO11-cls-m model for activity classification and the YOLO11m model for human detection. The collected dataset was divided into training, validation, and test sets with a 70%, 15%, and 15% split, respectively. The YOLO11-cls-m model was trained on 13 action categories using a YOLO model pre-trained on the COCO dataset, while YOLO11m was trained using a pre-trained model from ImageNet. To mitigate overfitting, we incorporated an early stopping mechanism and applied the default data augmentation techniques provided by YOLO11.

Result. The detailed classification results are presented in the Appendix. For our implementation, an activity was considered detectable if its F1 score exceeded 90%. We have summarized the activities that can be recognized, both with and without significant privacy concerns, at different resolutions in Tables 1. This information is used in our system to display a list of detectable activities at varying sensing resolutions. Most activities, particularly those with privacy concerns, were not detectable at the lowest resolution of 40×30. However, Walking and Falling Down remained detectable even at this resolution, allowing the system to track user movement and support the dynamic zone feature, which follows a user's position in real-time. We provide the sensor data in the Appendix to illustrate examples of user activities captured at different sensing resolutions. Note that our goal was to support informed decisions regarding which activities should be displayed as detectable by our system. Factors such as model architecture, training strategies, and data augmentation methods may result in different conclusions than those presented here. Therefore, readers should not interpret these results as a definitive reference for the detectability of activities at specific sensing resolutions.

6 OVERVIEW OF USER STUDIES

We conducted two user studies to evaluate the usability and effectiveness of our approach. The first study assessed how well the proposed visual feedback communicated critical information about sub-sensing zones during user interactions. Overall, participants found the visual feedback to be intuitive and easy to understand. This helped isolate potential confounding factors that could otherwise influence the privacy-related experiences examined in the second study. In the second study, we explored how effectively LiDArgus supported diverse privacy preferences. Findings indicate that LiDArgus improved participants' awareness of the LiDAR sensor's operational status and enhanced the transparency of its capabilities. It also offered effective mechanisms for privacy protection and user control over the baseline LiDAR system. Compared to a baseline condition without sub-zoning, the sub-zoning feature more effectively supported privacy preferences for both visitors and the LiDAR system's owner.

7 USER STUDY 1: UNDERSTANDING VISUAL FEEDBACK

The goal of this study was to assess the clarity of the visual feedback and its effectiveness in communicating information about sub-sensing zones. Specifically, we examined how well the feedback communicated key details, including the zones' location, spatial coverage, sensing resolution, and the types of activities detectable within each zone.

7.1 Study Design

This section describes the design of our study, including the participants, apparatus, tasks, and overall procedure.

Participants. We recruited 12 participants (8 male, 4 female), with a mean age of 25.5 (SD = 2.53). All had prior experience with voice input systems for smart home appliances, mobile devices, or vehicles. Moreover, all participants had normal vision and hearing.

Apparatus. We conducted the study using the prototype described in the Implementation section. The test environment and zones were the same as those described in the walkthrough. Participants used a Sony SRS-XE300 speaker for audio feedback and HyperX SoloCast for voice input.

Task. Participants used voice commands to trigger and evaluate visual feedback projected onto the room to assess its effectiveness in conveying the intended information. Below, we outline the tested feedback. For feedback triggered by voice commands, we also include the corresponding command (in brackets).

Create a new zone ("create zone"): Visual feedback showed the zone boundaries, raw data, and detectable activities (Figure 9).



Figure 6: (a) Alice initiates the exploration of Bob's room by issuing the voice command "follow me," which transitions her zone from static to dynamic mode, enabling continuous tracking of her movements. (b-c) This transition allows the zone to adjust in real-time to her location. The zone acts as a mask, preventing Alice from being detected at higher resolutions, even when she enters the high-resolution zone near Bob's desk.



Figure 7: In another scenario, where (a) Bob sits in the room on his rock chair. (b) Alice's zone is adjacent to Bob's, causing the two zones to merge. (c) The merged zone is now larger, covering both Alice and Bob. The system automatically adjusts the sensing resolution to the lower of the two zones, prioritizing the more restrictive privacy settings.



Figure 8: The hardware setup of LiDArgus, including the M30 LiDAR Development Kit from Lumotive, along with a projector.

Remove a zone ("remove zone"): The boundaries of the zone encompassing the participant were initially displayed and then gradually faded away (Figure 10).

Disable sensing ("disable sensing"): Visual feedback only showed the boundaries of the zone (Figure 11).

Adjust zone resolution ("increase/reduce resolution", "maximize/minimize resolution"): Visual feedback displayed the zone boundaries with the updated sensor data and corresponding activities that could be detected (Figure 12, Figure 13).

Toggle between dynamic and static mode ("follow me"/"stop following"): The dynamic mode was indicated by an icon displayed near the top of the zone. Participants were encouraged to walk across the room to observe how the dynamic zone adjusted to follow their movements.

User entering a zone: Visual feedback showed the zone boundaries, raw data, and detectable activities.

User entering the room: Visual feedback showed all zones in the room. Participants were encouraged to walk across the room to see how their body movements were sensed under varying sensing resolutions in each zone.



Figure 9: The visual feedback for the command "Create Zone".

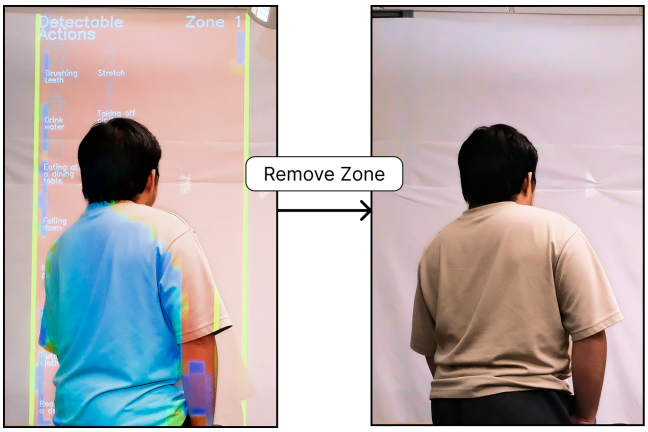


Figure 10: The visual feedback for the command "Remove Zone".

Two zones merging: Visual feedback displayed the merged larger zone with the resolution set to the lower of the two zones.

Participants were given as much time as needed to complete each task. Before each trial, the visual feedback was disabled, and it was turned off again after each trial, prior to starting the next one.

Procedure.

The tasks were presented to participants in a random order. Before the study, participants were introduced to the concept of zones and sensing resolution. However, they were not provided the opportunity to familiarize themselves with the tested visual feedback. This was intended to assess the usability of our design for first-time users, such as visitors in a smart home or office environment. Additionally, participants were not informed about the system's privacy-related goals, as the study aimed specifically to assess the clarity and usability of visual feedback for sub-zones, independent of any privacy implications.

7.2 Result.

For each tested visual feedback, participants shared their understanding of the message it aimed to convey. After receiving their response, we provided the correct answer. Subsequently, participants rated the following two statements on a 5-point Likert scale (1: strongly disagree and 5: strongly agree): (1) I immediately understood this feedback, and (2) the visual feedback clearly communicated the information the system intended to convey.

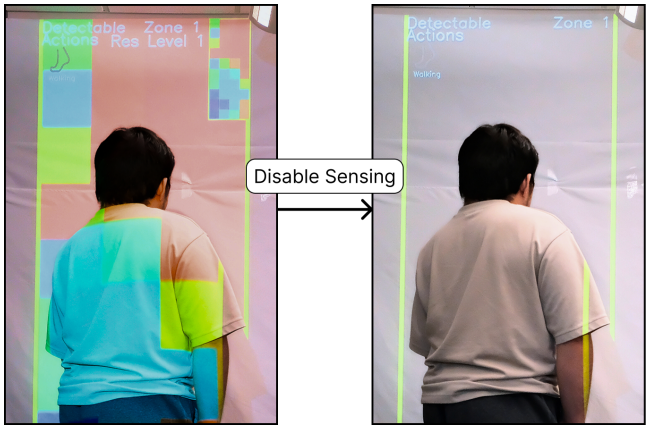


Figure 11: The visual feedback for the command "Disable Sensing".

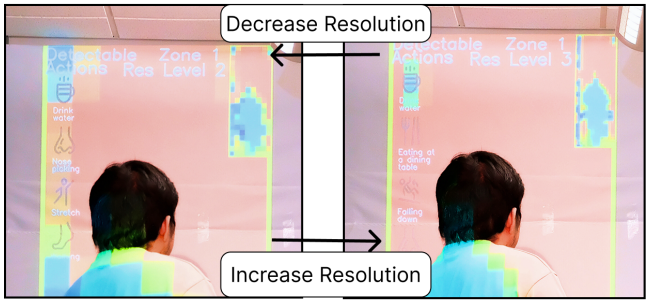


Figure 12: The visual feedback for commands "increase/decrease resolution".

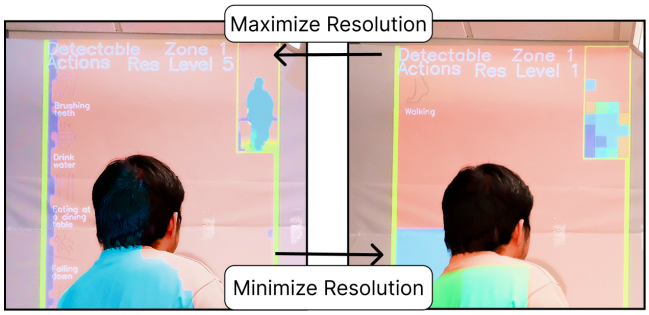


Figure 13: The visual feedback for command "minimize/maximize resolution".

Figure 14 presents the results of our study. Overall, participants reported that they were able to quickly understand the tested visual feedback and agreed that it effectively conveyed the intended system information. Among the tested visual cues, the feedback for "Remove a zone" received the highest consensus, with all participants rating it 5 for both ease of understanding and clarity of communication. For "User entering a zone", approximately 60% of participants strongly agreed that the feedback was both easy to understand and clearly communicated its intended message. Responses to the feedback for "Create a new zone" were more varied. Over 50% of participants reported that they did not immediately grasp its meaning. However, once they understood it, more than 70% agreed that it effectively conveyed the system's intended information. For the remaining visual feedback cues, over 70% of participants strongly agreed that they were both easy to understand and clearly communicated their intended message. Next, we discuss our findings and observations in more detail. It is important to note that participants' feedback in this study was limited to visual feedback cues and did not address their privacy preferences.

7.2.1 *Visual Feedback.*

Overall, participants agreed that the tested visual feedback effectively conveyed key operational information about the zones, including their boundaries, the raw data captured, the range of supported sensing resolutions, and the types of activities recognizable at each resolution level within the system's current implementation.

Raw sensor data. Participants found that displaying the raw sensor data was particularly necessary and informative for understanding the LiDAR's capabilities. As P3 noted, *"Since I am not very familiar with LiDAR, I really wanted to know how the sensor scans my body, and this feedback was very helpful. Without it, I might have been completely in the dark"*. Participants with prior exposure to LiDAR-generated raw point cloud data in a 3D environment also found the visual feedback intuitive when applied to this new context, where 2D pixels were projected onto their bodies. For example, P9 stated, *"The visualization of the raw sensor data was easy for me to understand the different sensing resolutions. I could clearly see the differences based on the size of the pixels outlining my body"*.

Color mapping. In our visualization, the depth of each pixel was mapped to an 8-bit color scale. While participants observed dynamic changes in the pixel colors, some found this representation initially unfamiliar (P4, P5, P7). P7 said *"I didn't quite understand what these differently colored blocks represent and why the color changed"*. Similarly, P4 noted, *"I eventually grasped the meaning of the color mapping after interacting with the system for a bit, but it still required some thought and guessing. It's not something I imagine people would understand immediately"*. However, these observations do not suggest that the visualization approach is ineffective. As P5 commented, *"At first, I didn't realize that these colors represented different distances from the sensor, but once I understood it, I found the colors to be a very intuitive way to convey depth"*.

Information overload. Presenting a large amount of information simultaneously left some participants feeling overwhelmed (P1, P2, P10, P12). In particular, some participants were immediately drawn to the raw sensor data, causing them to overlook the icons representing recognizable actions and the zone boundaries. As P12 noted, *"The scanning results caught all my attention, so much so that I didn't even notice there was an icon there"*. This finding, while not entirely unexpected, highlights an important design consideration: simply displaying all available information about sensing zones does not guarantee effective communication. Instead, the design of the output must also account for users' ability to process and interpret this information effectively.

On-body projection. Some participants (P8, P10, P11) noted that initially, seeing the raw sensor data projected directly onto their bodies caused some confusion. However, after moving around, they eventually noticed the light on their bodies and recognized the sensor pixels. For example, P10 said, *"When I didn't realize that the projected image was shown on my body, I thought my body movements were not being sensed. But then I realized that was not the case, and I'm glad I figured it out"*. This insight highlights the importance of ensuring that participants are aware of where information is displayed. While our intent was to convey the operation of the LiDAR zone, the lack of immediate recognition of the projection medium led to an initial misinterpretation. Some participants assumed the sensor had failed to detect them. One potential factor contributing to this misunderstanding was the use of different media to convey distinct types of information, for example, raw sensor data projected onto the body, and supplementary information displayed in the physical environment. This separation may have encouraged participants to form an incomplete or inaccurate mental model, expecting all relevant information to be accessible via a single medium. In addition to the cognitive demands of interpreting multiple displays, this setup also introduced physical overhead. Participants were required to shift their attention between their bodies and the surrounding space. As P1 pointed out, *"Although projecting onto the body is more intuitive, having to repeatedly turn around to view both the projection on my body and other information displayed on the wall was somewhat tiring. It would be much better if I could view all the information in one place"*.

7.2.2 *User Reactions to Visual Feedback.*

Here, we report participants' reactions to interacting with different forms of visual feedback.

Recognizable activities. Several participants expressed surprise when the list of recognizable actions changed as they toggled between different sensing resolutions (P5, P7, P9–P12). This observation highlighted that, while participants quickly understood the visual feedback, a more compelling outcome was how this feature encouraged them to experiment with the system's recognition capabilities (a behavior also observed with other forms of feedback described later). Many participants spontaneously engaged in actions not listed, testing the system's boundaries and observing its responses. These interactions seemed to facilitate a deeper understanding of the system's capabilities. For instance, P12 commented, *"At this resolution (*40×30*), even my arms are hard to see, so I guess even if I did something weird, the system probably wouldn't detect it"*. Similarly, P9 noted that *"It feels like this resolution (*40×30*) is too low for the system to recognize what I'm doing, so I guess it's safe?"*. Interestingly, participants' responses regarding privacy were unexpected, as they were not informed in advance that the feedback designs had a privacy-related focus. This spontaneous exploration

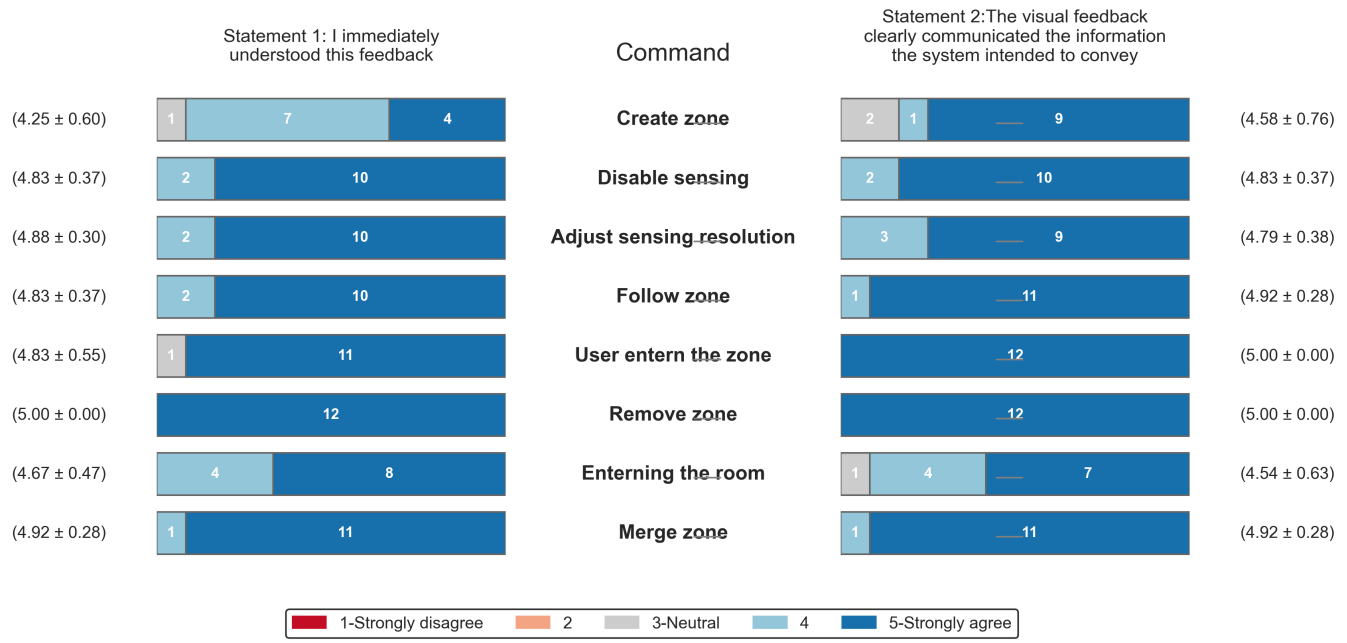


Figure 14: Participants' ratings on two questions regarding the tested visual feedback: (1) 'I immediately understood this feedback' and (2) 'The visual feedback clearly communicated the information the system was intended to convey'.

was especially encouraging, as it demonstrated the system's discoverability and how it could help users better understand potential privacy implications.

Sensing resolution. Participants were free to adjust the sensing resolution as often as they wished and to explore the full range of resolution levels. Notably, all participants eventually engaged with all ten resolution levels supported by our system. When asked about their motivations, many highlighted the satisfaction of being able to control and fine-tune the sensor's resolution, describing it as empowering and engaging. As P12 described, *"It is super satisfying to be able to adjust the resolution and see the sensing area shift, from such a high resolution, where my body contour is very sharp, to a much lower resolution, where it appears coarse"*. We view this as a compelling strength of our design, as it invites active exploration and fosters a sense of agency. This type of interaction may also encourage users to think more deeply about privacy considerations, increasing their awareness of the system's implications. While the overall response was positive, a few participants (e.g., P2, P6) observed that differences between adjacent resolution levels were sometimes too subtle to be easily distinguishable. As P6 explained, *"The differences between adjacent resolution levels are subtle, which makes me uncertain which level is useful for what and how it would affect my experience"*. P2 echoed this sentiment, suggesting that *"For someone who does not own this system, such fine-grained resolution levels might not be necessary. A simpler scheme with low, medium, and high settings could be more intuitive"*. These reflections indicate that participants not only grasped the system's visual feedback but also thoughtfully considered its usability and suggested avenues for refinement.

Dynamic zone. When the new zone was switched to dynamic mode, several participants (P2–P7, P10–P12) initially appeared uncertain. They approached the zone cautiously, edging toward its boundary, and only began walking at a normal pace once they noticed that the zone actively responded to their movements by updating its location. As participants became familiar with how the dynamic zone operated, many (P2, P4, P6, P10–P12) began moving across the room at varying speeds, engaging in exploratory behavior to better understand and test its capabilities. This pattern suggests a willingness to interact with and learn from the system, even in the face of initial ambiguity. P9 described the experience as *"I felt as if I were under a spotlight that follows me,"* highlighting how the dynamic zone evoked a sense of being recognized or tracked. While this sensation was intriguing for many, some (P1, P2) noted that it felt slightly awkward when they were uncertain about the system's sensing capabilities. For instance, P1 shared, *"I felt a bit uneasy at the beginning until I paid attention to the raw sensor data. Then it was fine. It was an experience I hadn't had before, and perhaps I will eventually get used to it."*

Fixed zone. Participants engaged with the boundaries between adjacent sensing zones in diverse and thoughtful ways. Many (P1–P5, P7, P9–P12) positioned themselves at the boundary between high- and low-resolution zones to observe how each side of their body was scanned differently. Several participants (P2, P5, P6, P7, P9, P10) actively stepped in and out of zones to explore how the sensing resolution shifted across their body. Others (P3–P5, P10–P12) extended their arms across zone boundaries to experience the contrast in resolution. P5 said, *"It is super cool to have the sensor only detect my hands and be able to see it"*, highlighting the novelty of the experience. Similarly, P7 noted the unique quality of the interaction

compared to previous experiences with LiDAR or camera-based systems: *"Having my arm sensed at two different resolutions made it feel as though my arm was divided into two parts"*. Participants adapted quickly to these interactions and found them helpful for understanding the distinctions between zones. Most were able to correctly identify all three zones by observing the system's visual output. However, a few (P1, P2, P5, P9, P12) initially overlooked the zone where sensing was disabled, as only the zone boundaries were displayed for this area.

7.2.3 Improvements.

Based on the findings from this study, we implemented three design improvements before proceeding with the second user study. First, to address the issue of raw sensor data being displayed separately from other information, we added a thumbnail of the sensor data in the upper right corner of each zone. This adjustment ensures that users are not restricted to their body surface when viewing the complete raw sensor data. Second, to better inform users about zones with lower salience (e.g., zones where sensing is disabled), we introduced color coding for the interior of these zones. When a user enters such a zone, the color appears and gradually fades over two minutes, visually indicating that no sensor data is being captured within that area. Third, we improve the clarity of the interface by adding text descriptions to activity icons, making them both more recognizable and descriptive. Additionally, we numbered the zones and included text labels (e.g., "Zone 1") to improve differentiation between them.

8 USER STUDY 2: THE EFFECTIVENESS OF LIDARGUS

The goal of this study was to understand the user experience of LiDArgus. Specifically, we conducted interviews with participants to understand how our approach aligns with their privacy preferences in the roles of both space owner and visitor. As space owners, participants performed tasks such as working at a desk and relaxing on a rock chair. As visitors, they were asked to explore the space, simulating the experience of a prospective tenant assessing a room for potential rental. In addition to LiDArgus, we included a baseline condition where the LiDAR operated without sub-sensing zones.

8.1 Study Design

This section presents our study design.

Participants. To ensure that participants were familiar with our visual feedback system, we selected the same group of 12 participants from Study 1. All participants reported having prior experience with activity-sensing devices, such as smart motion sensors or security cameras. The study was conducted one week after Study 1 to mitigate potential short-term memory effects that might influence participants' perceptions of our approach.

Apparatus. We conducted the study using the prototype described in the Implementation section. In the LiDArgus condition, the space was divided into three zones, as described in the walk-through. In the baseline condition, the entire space was covered by a single zone with the highest sensing resolution, representing the typical use of LiDAR for activity sensing. No visual feedback was provided in this condition. Because many work- and health-related activities, such as reading or drinking water, cannot be reliably recognized at lower resolutions (see Figures 19 and 20 in the appendix), we set the maximum sensing resolution to 640×480 in both conditions.

Task. In both the LiDArgus and baseline conditions, participants completed tasks by following voice-guided, step-by-step instructions. Progression to the next step was triggered by the participant saying "next."

In the LiDArgus condition, role-playing as a visitor, participants were asked to imagine themselves as a prospective tenant evaluating a room for rent while the owner was temporarily absent. They were not informed whether the LiDAR sensor was active and were expected to determine this on their own. The task involved nine steps: (1) entering the room (all static zones appeared); (2) creating their own zone (new zone appeared, static zones invisible); (3) adjusting resolution settings based on privacy preferences (all zones visible); (4) switching the zone mode to dynamic, allowing it to follow their movement (all zones visible); (5) walking inside the room (all zones visible); (6) removing their jacket (all zones visible); (7) sitting at the desk and using their phone (all zones visible); (8) picking their nose (all zones visible); and (9) returning to the door and exiting (animation showed new zone removal). These steps encompassed interactions involving observation of all zones in the room and their sensing settings, control over their designated zone, navigation of the room under their chosen sensing setting, and engagement in actions they might prefer to remain undetected by the sensor and unseen by space owner at a later time, such as removing their jacket, using their phone at the owner's desk, and picking their nose.

In the LiDArgus condition, where participants role-played as space owners, they were asked to imagine that they had installed the LiDAR sensor to monitor their own work- and health-related activities. They followed a set of voice-guided instructions to complete a sequence of seven tasks: (1) stepping into the zone encompassing the desk (high-resolution zone briefly appeared and faded away after two minutes); (2) reading at the desk (no zone visible), (3) drinking water at the desk (no zone visible), (4) picking their nose at the desk (no zone visible), (5) walking to the rock chair (sensing-disabled zone briefly appeared and faded away after two minutes), (6) taking off their jacket (no zone visible), and (7) scratching their back (no zone visible). These steps included activities that the owner might prefer not to be recorded by the system, such as picking their nose, taking off their jacket, and scratching their back. They took place in areas that might or might not be detectable by the sensor, such as the desk and rock chair. Zone visualization was mostly disabled, as we anticipated that owners would find it unnecessary in daily use.

Procedure. Prior to the study, participants were introduced to LiDAR and the key features of LiDArgus, including the sensor's capabilities, sub-sensing zones, available controls, and feedback. They were also shown examples of raw sensor data at various resolutions along with corresponding machine learning classification results. Participants were encouraged to think aloud throughout the study. To prevent bias, participants first played as visitors in the baseline condition. They were not informed whether the LiDAR sensor was turned on or off, simulating real-world scenarios where visitors may be unaware of sensing activities. Subsequently, participants completed the remaining role-condition combinations in

a randomized order. Before their trial as an owner in either condition, participants were informed about the sensor's operation and given the opportunity to review its functionality, regardless of whether they had already encountered this information in previous LiDArgus conditions. This ensured that owners had a clear understanding of the system. During the study, participants were left alone in the room. Following the study, each participant took part in a semi-structured interview.

8.2 Result

Participants responded to a series of questions addressing their awareness of the sensor's operational status and transparency, the availability of privacy protection options, and how well the system supported them in expressing their privacy preferences. Additionally, they were asked to evaluate whether, and to what extent, the visualizations provided by our approach were helpful in supporting privacy-related decision-making. These visualizations included features such as the ability to view zone boundaries, inspect sensing resolution levels, access raw sensor data, and understand which activities were detectable. Participants rated their responses using 5-point Likert scale items (1 = Strongly Disagree, 5 = Strongly Agree). Statistical analysis was performed using the Friedman test.

Figure 15 and Figure 16 present the results from the perspectives of the owner and the visitor, respectively. Overall, participants expressed strong agreement that our approach effectively provides meaningful information and control, addressing several key aspects of their privacy preferences.

8.2.1 User Awareness of LiDAR Operational Status and System Transparency.

LiDArgus scored higher in both user awareness of LiDAR operation (Owner: $X^2 = 2.42p < 0.05$, Visitor: $X^2 = 3.11, p < 0.01$) and transparency regarding the system's functionality and capabilities (Owner: $X^2 = 2.59, p < 0.01$, Visitor: $X^2 = 3.16, p < 0.01$). Most participants reported that the information provided by LiDArgus helped them better understand the operational status of the LiDAR system, the coverage of its sub-sensing zones, and its sensing capabilities. This increased understanding helped alleviate their privacy concerns. User perspective of the system had no significant difference when participants were in different roles ($X^2 = -3.09, p < 0.01$).

When experiencing the baseline system as visitors, participants often felt uncertain about when, where, and how their body movements were being sensed. As P12 said, *"I wasn't sure if the sensor was on or off, and I felt like I had no control—this was awkward"*. Similarly, P9 shared, *"I thought about covering the sensor with a cloth before doing weird things like picking my nose, but in reality, that would be in someone else's house, and I'm not sure if I'm allowed to do that"*. Echoing findings from prior work [26, 41, 74], such ambiguity can lead individuals to adapt or compromise their privacy preferences. Interestingly, participants responded more positively to the baseline condition when adopting the role of the owner, as familiarity with the system's operation provided reassurance. While visual feedback about the sensor's status was not deemed essential (33% participants), many participants expressed a desire for it. For instance, P9 and P11 noted, *"As an owner, I'm more concerned about whether the system is functioning properly. Right now, I don't know if it's operating according to the parameters that I set. I'm not sure whether the LiDAR is correctly monitoring my activities"*. However, this confidence often shifted when participants assumed the role of a visitor. As P9 reflected, *"After I entered the room, I did not notice where the sensor was located, nor did I understand how it worked. This sense of the unknown was a bit scary"*.

In contrast, under the LiDArgus condition, participants (P1, P2, P4, P6, P10, P11) reported that the feedback helped them intuitively recognize when the LiDAR sensor was actively sensing them. Even when participants expressed initial discomfort about being sensed, they appreciated the system's transparency. For example, P11 said that *"When I saw the visualized information, I immediately knew there was a sensor observing me and felt an unsettling sense of surveillance, yet the visual feedback helped me understand what kind of data was being captured, which made me feel better"*. Similarly, P2 shared, *"Although it's somewhat awkward to see myself being sensed, LiDArgus informs me of different zones. I could choose to stay within a zone with lower sensing resolution—something the baseline didn't offer"*. Participants expressed greater enthusiasm when acting as visitors compared to when taking the role of owners. While the volume of information initially required some effort to interpret, participants, especially in the visitor role, were generally willing to explore different sensing resolutions before deciding on their next action. These findings suggest that making the operation of the LiDAR system transparent can significantly enhance users' comfort and willingness to engage with the technology.

8.2.2 Privacy Protection Options and Expression of Privacy Preferences.

LiDArgus received higher scores for the availability of privacy protection options (Owner: $X^2 = 2.88, p < 0.01$, Visitor: $X^2 = 3.10, p < 0.01$) and the ability of participants to freely express their privacy preferences compared to the baseline condition (Owner: $X^2 = 2.74, p < 0.01$, Visitor: $X^2 = 3.04, p < 0.01$).

In the baseline condition, participants expressed notable concerns regarding the limited control over the sensing system. While many participants (P3–P6, P10–P12) recognized that, as owners, they could manually toggle the sensor on or off, they felt this binary control was insufficient to accommodate their diverse privacy preferences. For instance, P6 stated, *"I can only control whether it is on or off, and this means I would have to toggle the device frequently if I just want to temporarily disable it. The worst part is, I doubt I could remember to turn it off every time I plan to do something awkward that I wouldn't want to be sensed. It's not easy"*. When placed in the role of a visitor, several participants (P11, P6, P4, P3) described a sense of discomfort stemming from the sensor's persistent monitoring, especially when they lacked meaningful agency over how their data was handled. As P4 shared, *"I have no way to do anything to change it. It is frustrating that I have no other option except accepting that I have no control over how my body is sensed"*. Importantly, participants' desire for control could be influenced by the length of their stay. P10 reflected, *"If I remain in the area for only a few minutes, I might not worry about not having control over the system, but if I have to stay longer, I would be more concerned that I might inadvertently engage in actions that compromise my privacy, like subconsciously picking my nose or scratching my back while watching TV"*.

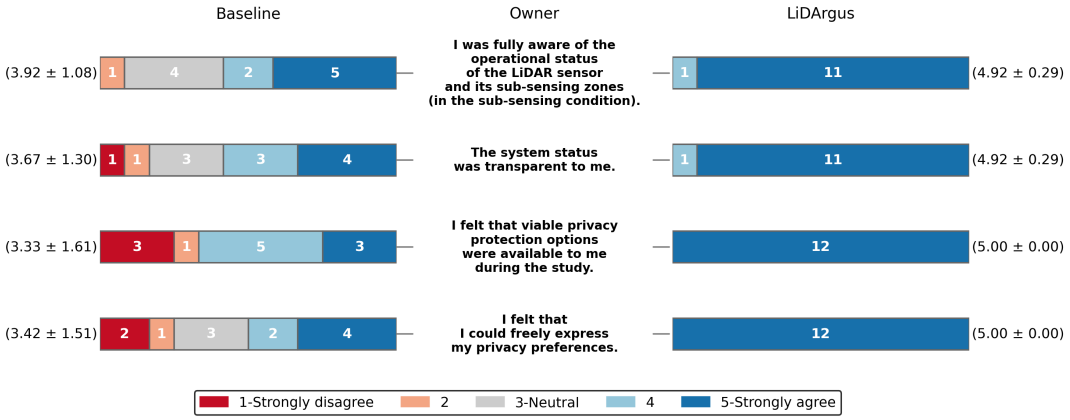


Figure 15: Likert ratings of participants for the baseline condition and the LiDArgus condition while assuming the role of space owner.

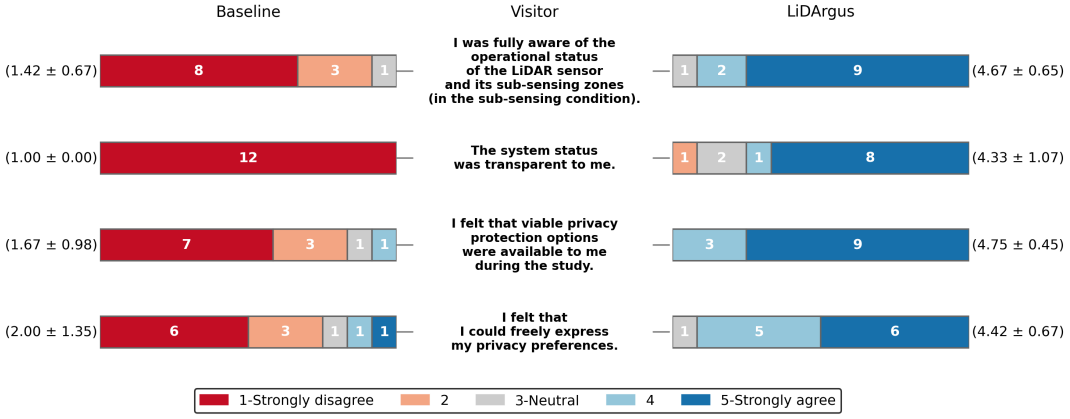


Figure 16: Likert ratings of participants for the baseline condition and the LiDArgus condition while assuming the role of visitor.

Participants appreciated the enhanced sense of control afforded by the LiDArgus condition when taking on the role of the homeowner. They valued that LiDArgus provided more nuanced options beyond merely toggling the LiDAR sensor on or off, allowing them to tailor privacy settings to their individual preferences. The ability to manage sub-sensing zones was especially well-received. For example, P1, P3, P9, and P11 expressed that LiDArgus enabled them to feel significantly more in control of the sensor, as opposed to using ad-hoc methods such as covering it with a cloth when desiring temporary deactivation. As P9 shared, *"It is hard to imagine anything else I would want to control. As long as the data is stored locally, I am comfortable using the LiDAR sensor at home"*.

When the role shifted to that of a visitor, participants expressed appreciation for the ability to create and modify sub-sensing zones, rather than being limited to the default zones predefined by the owner. They regarded having control over the LiDAR sensor as even more valuable than merely being shown the sensor's operational status. Interestingly, after creating a new zone in the visitor role, some participants explored different sensing resolution levels and chose settings higher than the default minimum of 40×30. For example, P12 selected a resolution of 160×120. When asked about their choice, P12 shared that privacy was less of a concern at that moment, making a slightly higher resolution acceptable. As P12 explained, *"I now understand the capability of the system and how it works, and I feel I can properly evaluate the privacy implications of different sensing resolutions, so I thought I'd be okay using a higher resolution than the default"*. While prior research has noted that users can sometimes make suboptimal privacy decisions [3, 11, 31, 66], our findings highlight a promising outcome: users, especially visitors, were able to make informed decisions that aligned with their own privacy preferences, suggesting increased awareness and agency in managing privacy settings.

Conversely, participants raised thoughtful concerns about scenarios involving multiple users. P2 highlighted a potential privacy issue, noting, *"If sensor data is shown in a scenario with multiple people, even a small awkward movement I make might be visible to everyone. That would be super embarrassing"*. Usability in social contexts also emerged as a consideration. Reflecting on a home party scenario, P2 shared, *"If many guests at a party are visiting for the first time, they would need to familiarize themselves with the system and take time to learn how to use it. I worry that a lot of people won't have the patience to do it properly. And as the owner, I definitely can't explain it to everyone. Training them all at once would be impossible"*. These concerns are understandable, especially given that the system was entirely new to participants. However, such challenges may diminish as the technology becomes more commonplace. As P4 optimistically noted, *"As more people use the system, I believe that training time will no longer be necessary"*.

8.2.3 LiDArgus Features.

Next, we present participants' reactions to the concept of sub-zoning and the visualized information provided by LiDArgus. Note that the insights gained in this part of the study focus on privacy preferences, which were not addressed in Study 1. The results are shown in Figure 17.

Sub-sensing Zone. Participants generally found the concept of sub-sensing zones effective in meeting their privacy needs, as previously discussed. Notably, the dynamic zone, which adapts to the user's movement, appeared to be a preferred feature. Participants appreciated using it as a "mark" to navigate within fixed zones offering higher sensing resolutions, without concern for their privacy-sensitive activities being detected. P4 shared, *"I like this zone. It follows me well. Otherwise, I would have to go to a specific low-resolution zone or one where sensing is disabled to feel safe"*. P10 further commented, *"I feel like I own this place, even while playing the role of a visitor. I hide in my shield, it follows me, and I know my activity won't be sensed by the owner's high-resolution system"*. P3 expressed, *"Although I'm not familiar with how the system implements that mask, I'm confident that it actively safeguards my privacy, which is very reassuring"*.

Zone Boundaries. Participants generally found the visualization of zone boundaries to be a valuable aid in making decisions aligned with their privacy preferences. Most participants in the owner role felt that, since they were responsible for configuring the system, it was unnecessary to display zone boundaries at all times. However, they appreciated being able to glance at the boundaries upon entering a zone, a behavior that our system supports. As P2 noted, *"I liked that when I entered the room, I could see the zones, which also lets me know if the system is operating properly"*. In contrast, all visitor-role participants emphasized the importance of clearly displaying zones and their boundaries. In addition to 2D projections of the boundaries, several participants suggested visualizing sub-zone boundaries in 3D space. For instance, P1 said, *"Since the LiDAR scans the area in a fan shape, it would be great to see that fan in 3D"*. We believe such visualizations could be effectively supported using AR glasses.

Sensing Resolution. All participants emphasized the importance of displaying the sensing resolution. For owners, it is important for making informed decisions when subdividing space. For visitors, the sensing resolution serves as a clear indicator of the system's current sensing capabilities, thereby aiding in understanding the privacy implications.

Raw Sensing Data. Participants' attitudes toward raw sensing data varied based on their role. For owners, some participants (P6, P7) did not prioritize the raw data. As P6 explained, *"Once I have set up the system, there is no need for the raw data to be displayed unless I need to modify a specific sensing zone"*. P7 shared a similar view, stating, "As an owner, I guess there is no need for me to see the raw data all the time". In contrast, visitors placed more importance on the raw data. Beyond the 2D representation, some participants expressed interest in viewing the raw point cloud data in 3D. For instance, P7 commented, *"Currently, the visual feedback only shows the data in 2D, but this sensor (LiDAR) actually scans in 3D. It would be better if it could show the exact 3D scanning data so I won't miss anything"*.

Detectable Activities. As a key feature of the system, most participants in the owner role found it beneficial to know which actions could be detected at various resolutions. As P8 noted, *"This clearly shows that the system is sensing and predicting my actions, letting me know that it has seen me"*. However, some participants in the visitor role (P1, P6) expressed that, instead of focusing on the specific actions detected, they were more concerned with a simpler question: whether the LiDAR sensor was configured to detect sensitive activities—yes or no. If the answer was yes, they were less focused on the precise activities detected and more interested in being informed and cautious. For instance, P1 mentioned, *"I only care whether the LiDAR sensor is configured to capture sensitive actions. To me, it seems that the list of detectable activities is designed for owners"*.

9 LIMITATION AND FUTUREWORK

In this section, we discuss the limitations of our approach, reflect on the lessons learned throughout the course of this research, and outline directions for future work.

Sub-Zoning. We propose the concept of sub-zoning as a way to better accommodate the diverse privacy preferences of different user groups simultaneously. In this approach, privacy is prioritized over other factors, such as a user's preference for system functionality. For example, when a visitor's zone overlaps with the owner's zone, the resulting area is automatically assigned a lower sensing resolution to safeguard the visitor's privacy. This, however, comes with the trade-off of reduced activity detection for the owner. Future research should investigate alternative strategies that can more effectively balance varying user preferences, including privacy and functionality. This area is still relatively unexplored and offers a promising direction for further research.

1D Sub-Zoning. We demonstrate the effectiveness of sub-zoning by partitioning the LiDAR's sensing region along the x-axis. This approach is driven by current hardware limitations, which support one-dimensional segmentation of the field of view. While this 1D partitioning effectively illustrates the sub-zoning concept in its initial stages, it also highlights opportunities for further development. Specifically, the current implementation restricts the ability to selectively sense specific body parts, such as the hands, with high resolution while masking others, like the face. Additionally, it does

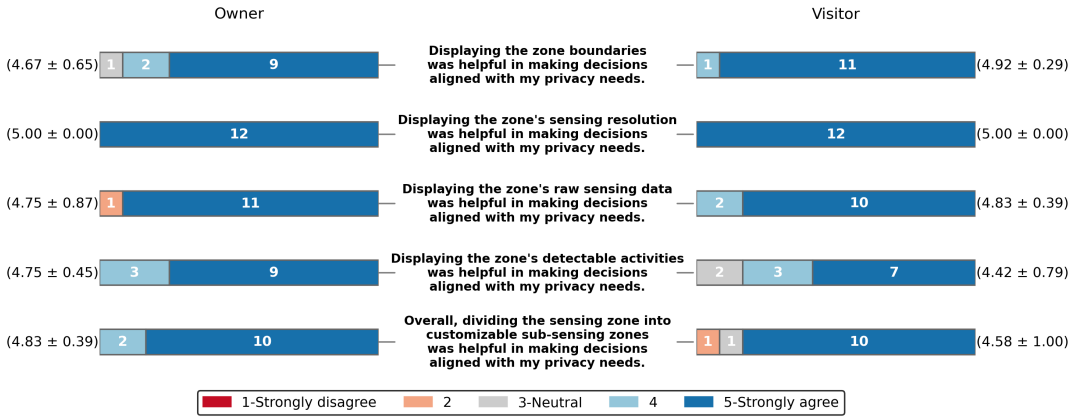


Figure 17: Likert scale ratings provided by participants for the different features of LiDArgus.

not allow zones to conform to the contours of objects, such as a couch. Future work should explore new technical advancements that enable more flexible, multidimensional zoning and examine their potential to better accommodate diverse user privacy preferences.

Shared Visual Feedback. The current implementation of LiDArgus offers shared visual feedback that is visible to all users co-located in the same environment. This approach is simple to implement and serves as an effective solution for early-stage exploration. However, it also introduces potential privacy concerns, as a user's activity may be publicly displayed in the physical space, potentially exposing it to nearby individuals. Future research could explore the benefits and limitations of this projection-based visualization approach and investigate alternative methods, such as the use of AR glasses, to enhance privacy while maintaining user engagement.

Visualizing Zone Boundaries in 2D. In our current implementation, zone boundaries are projected onto the physical environment and visualized in 2D. While this method effectively communicates the location and boundaries of the zones, it presents certain limitations. Specifically, LiDAR sensors scan the environment in 3D, yet the zone boundaries are represented in 2D, which may hinder users' ability to fully perceive the zones' coverage. In future iterations, we plan to explore alternative display devices, such as AR glasses, to visualize the zones in 3D. This enhancement will allow us to assess its impact on the user experience and uncover insights that may not have been apparent in the current study.

Projector. Our current implementation utilizes a projector to provide visual feedback. While functional, this setup is cumbersome and less practical due to the requirement for pairing a LiDAR system with the projector. Future research should investigate alternative solutions that enable users to access information about sub-sensing zones, ideally by leveraging devices they already own. We aim to explore ways that users can harness their smartphones, whether through built-in cameras or simple add-ons, such as laser diodes, to detect beam-scanning regions and patterns. This would allow users to reconstruct zone boundaries and sensing resolutions, empowering them to independently detect LiDAR zones and eliminating the need for external output devices from manufacturers to address privacy concerns.

Study Design and Environment. Our user experiments were conducted in a controlled environment, which, although somewhat limited in external validity, allowed us to uncover several key insights regarding the effectiveness of our approach and highlight potential areas for future improvement. While real-world usage scenarios may present additional complexity and differ from the controlled setting, it is crucial to validate the concept of sub-zoning in more realistic environments. Both user studies involved the same group of participants. Although the studies were conducted one week apart to mitigate potential priming effects, we acknowledge that such effects cannot be entirely ruled out. This limitation may introduce bias toward either condition. To address this, future studies will recruit a broader participant pool to validate our findings. These efforts will help assess the broader applicability of our design and determine which findings remain robust across contexts, offering valuable guidance for refining and extending our approach.

User Misunderstandings of Privacy Protections. While reducing resolution can help mitigate privacy concerns in many cases, it does not fully guarantee privacy. In extreme scenarios, even very low-resolution data can result in privacy leakage. From a user perspective, this presents a risk of misunderstanding, specifically, the assumption that lower resolution equates to complete privacy, when in fact it only provides improved protection in certain scenarios. Future work should explore strategies for more effectively communicating these limitations to users. Another important direction for future research is the design of sensors with minimal risks of oversensing, for example, sensors that capture only task-specific data. We envision that combining efforts in interface design (as explored in our current work) with sensor-level innovations will better support user privacy preferences moving forward.

10 CONCLUSION

In this paper, we present a novel approach for accommodating diverse privacy preferences in LiDAR-based activity recognition

systems. Our method partitions the LiDAR sensing region into configurable sub-zones, each customizable in terms of location and sensing capabilities to reflect the privacy needs of different user groups. We validate this approach through a proof-of-concept prototype built with Lumotive's M30 LiDAR Development Kit, alongside two user studies. These studies evaluate the usability of our visual feedback mechanisms and demonstrate the effectiveness of our design in enhancing system transparency and user control. Importantly, our findings show that users in different roles can independently adjust their privacy settings without affecting others. This work represents an important step toward better balancing privacy and functionality in shared sensing environments. The insights gained from our studies can inform future research on user-centered privacy design in LiDAR systems and broader ubiquitous computing contexts.

References

[1] [n. d.]. Google Speech-to-Text API. https://cloud.google.com/speech-to-text/docs. Accessed: 2025-04-9.

[2] [n. d.]. M30 LiDAR Development Kit from Lumotive. https://lumotive.com. Accessed: 2025-04-9.

[3] Alessandro Acquisti, Laura Brandimarte, and George Loewenstein. 2015. Privacy and human behavior in the age of information. *Science* 347, 6221 (2015), 509–514. doi:10.1126/science.aaa1465 arXiv:https://www.science.org/doi/pdf/10.1126/science.aaa1465

[4] Amira Ben Mabrouk and Ezzeddine Zagrouba. 2018. Abnormal behavior recognition for intelligent video surveillance systems: A review. *Expert Systems with Applications* 91 (2018), 480–491. doi:10.1016/j.eswa.2017.09.029

[5] Carlos Bermejo Fernandez, Petteri Nurmi, and Pan Hui. 2021. Seeing is Believing? Effects of Visualization on Smart Device Privacy Perceptions. In *Proceedings of the 29th ACM International Conference on Multimedia* (Virtual Event, China) *(MM '21)*. Association for Computing Machinery, New York, NY, USA, 4183–4192. doi:10.1145/3474085.3475552

[6] Luca Catarinucci, Danilo de Donno, Luca Mainetti, Luca Palano, Luigi Patrono, Maria Laura Stefanizzi, and Luciano Tarricone. 2015. An IoT-Aware Architecture for Smart Healthcare Systems. *IEEE Internet of Things Journal* 2, 6 (2015), 515–526. doi:10.1109/JIOT.2015.2417684

[7] Jiawei Chen, Jonathan Wu, Kristi Richter, Janusz Konrad, and Prakash Ishwar. 2016. Estimating head pose orientation using extremely low resolution images. In *2016 IEEE Southwest Symposium on Image Analysis and Interpretation (SSIAI)*. 65–68. doi:10.1109/SSIAI.2016.7459176

[8] Wen-Tsuen Chen, Po-Yu Chen, Wei-Shun Lee, and Chi-Fu Huang. 2008. Design and Implementation of a Real Time Video Surveillance System with Wireless Sensor Networks. In *VTC Spring 2008 - IEEE Vehicular Technology Conference*. 218–222. doi:10.1109/VETECS.2008.57

[9] Emily Y. Chew, Traci E. Clemons, Susan B. Bressler, Michael J. Elman, Ronald P. Danis, Amitha Domalpally, Jeffrey S. Heier, Judy E. Kim, and Richard Garfinkel. 2014. Randomized Trial of a Home Monitoring System for Early Detection of Choroidal Neovascularization Home Monitoring of the Eye (HOME) Study. *Ophthalmology* 121, 2 (2014), 535–544. doi:10.1016/j.ophtha.2013.10.027

[10] Eun Kyoung Choe, Sunny Consolvo, Jaeyeon Jung, Beverly Harrison, and Julie A. Kientz. 2011. Living in a glass house: a survey of private moments in the home. In *Proceedings of the 13th International Conference on Ubiquitous Computing* (Beijing, China) *(UbiComp '11)*. Association for Computing Machinery, New York, NY, USA, 41–44. doi:10.1145/2030112.2030118

[11] Hanbyul Choi, Jonghwa Park, and Yoonhyuk Jung. 2018. The role of privacy fatigue in online privacy behavior. *Computers in Human Behavior* 81 (2018), 42–51. doi:10.1016/j.chb.2017.12.001

[12] Richard Chow. 2017. The Last Mile for IoT Privacy. *IEEE Security & Privacy* 15, 6 (2017), 73–76. doi:10.1109/MSP.2017.4251118

[13] Jessica Colnago, Yuanyuan Feng, Tharangini Palanivel, Sarah Pearman, Megan Ung, Alessandro Acquisti, Lorrie Faith Cranor, and Norman Sadeh. 2020. Informing the Design of a Personalized Privacy Assistant for the Internet of Things. In *Proceedings of the 2020 CHI Conference on Human Factors in Computing Systems* (Honolulu, HI, USA) *(CHI '20)*. Association for Computing Machinery, New York, NY, USA, 1–13. doi:10.1145/3313831.3376389

[14] Anupam Das, Martin Degeling, Daniel Smullen, and Norman Sadeh. 2018. Personalized Privacy Assistants for the Internet of Things: Providing Users with Notice and Choice. *IEEE Pervasive Computing* 17 (07 2018), 35–46. doi:10.1109/MPRV.2018.03367733

[15] R. Divya and J. Peter. 2021. Smart healthcare system-a brain-like computing approach for analyzing the performance of detectron2 and PoseNet models for anomalous action detection in aged people with movement impairments. *Complex & Intelligent Systems* 8 (03 2021). doi:10.1007/s40747-021-00319-8

[16] Pardis Emami-Naeini, Sruti Bhagavatula, Hana Habib, Martin Degeling, Lujo Bauer, Lorrie Faith Cranor, and Norman Sadeh. 2017. Privacy expectations and preferences in an IoT world. In *Proceedings of the Thirteenth USENIX Conference on Usable Privacy and Security* (Santa Clara, CA, USA) *(SOUPS '17)*. USENIX Association, USA, 399–412.

[17] Bohao Fan, Siqi Wang, Wenxuan Guo, Wenzhao Zheng, Jianjiang Feng, and Jie Zhou. 2023. Human-M3: A Multi-view Multi-modal Dataset for 3D Human Pose Estimation in Outdoor Scenes. arXiv:2308.00628 [cs.CV] https://arxiv.org/abs/2308.00628

[18] Paolo Fraccaro, Xenophon Evangelopoulos, and Blair Edwards. 2020. Development and Preliminary Evaluation of a Method for Passive, Privacy-Aware Home Care Monitoring Based on 2D LiDAR Data. In *Artificial Intelligence in Medicine: 18th International Conference on Artificial Intelligence in Medicine, AIME 2020, Minneapolis, MN, USA, August 25–28, 2020, Proceedings* (Minneapolis, MN, USA). Springer-Verlag, Berlin, Heidelberg, 160–169. doi:10.1007/978-3-030-59137-3_15

[19] Andrea Frome, German Cheung, Ahmad Abdulkader, Marco Zennaro, Bo Wu, Alessandro Bissacco, Hartwig Adam, Hartmut Neven, and Luc Vincent. 2009. Large-scale privacy protection in Google Street View. In *2009 IEEE 12th International Conference on Computer Vision*. 2373–2380. doi:10.1109/ICCV.2009.5459413

[20] A. Glandon, L. Vidyaratne, N. Sadeghzadehyazdi, Nibir K. Dhar, Jide O. Familoni, Scott T. Acton, and K. M. Iftekharuddin. 2019. 3D Skeleton Estimation and Human Identity Recognition Using Lidar Full Motion Video. In *2019 International Joint Conference on Neural Networks (IJCNN)*. 1–8. doi:10.1109/IJCNN.2019.8852370

[21] M. Greiffenhagen, D. Comaniciu, H. Niemann, and V. Ramesh. 2001. Design, analysis, and engineering of video monitoring systems: an approach and a case study. *Proc. IEEE* 89, 10 (2001), 1498–1517. doi:10.1109/5.959343

[22] Nawfal Guefrachi, Jian Shi, Hakim Ghazzai, and Ahmad Alsharoa. 2024. Leveraging 3D LiDAR Sensors to Enable Enhanced Urban Safety and Public Health: Pedestrian Monitoring and Abnormal Activity Detection. In *2024 46th Annual International Conference of the IEEE Engineering in Medicine and Biology Society (EMBC)*. 1–5. doi:10.1109/EMBC53108.2024.10782331

[23] Akshay Gupta, Shreyansh Jain, Pushpa Choudhary, and Manoranjan Parida. 2024. Dynamic object detection using sparse LiDAR data for autonomous machine driving and road safety applications. *Expert Systems with Applications* 255 (2024), 124636. doi:10.1016/j.eswa.2024.124636

[24] Xiao Han, Yiming Ren, Peishan Cong, Yujing Sun, Jingya Wang, Lan Xu, and Yuexin Ma. 2024. Gait Recognition in Large-scale Free Environment via Single LiDAR. In *Proceedings of the 32nd ACM International Conference on Multimedia* (Melbourne VIC, Australia) *(MM '24)*. Association for Computing Machinery, New York, NY, USA, 380–389. doi:10.1145/3664647.3681166

[25] Xiao Han, Yiming Ren, Yichen Yao, Yujing Sun, and Yuexin Ma. 2024. Towards Practical Human Motion Prediction with LiDAR Point Clouds. In *Proceedings of the 32nd ACM International Conference on Multimedia* (Melbourne VIC, Australia) *(MM '24)*. Association for Computing Machinery, New York, NY, USA, 7629–7638. doi:10.1145/3664647.3680720

[26] Julie Haney, Yasemin Acar, and Susanne Furman. 2021. "It's the Company, the Government, You and I": User Perceptions of Responsibility for Smart Home Privacy and Security. In *30th USENIX Security Symposium (USENIX Security 21)*. USENIX Association, 411–428. https://www.usenix.org/conference/usenixsecurity21/presentation/haney

[27] Roberto Hoyle, Robert Templeman, Steven Armes, Denise Anthony, David Crandall, and Apu Kapadia. 2014. Privacy behaviors of lifeloggers using wearable cameras. In *Proceedings of the 2014 ACM International Joint Conference on Pervasive and Ubiquitous Computing* (Seattle, Washington) *(UbiComp '14)*. Association for Computing Machinery, New York, NY, USA, 571–582. doi:10.1145/2632048.2632079

[28] Panagiotis Ilia, Iasonas Polakis, Elias Athanasopoulos, Federico Maggi, and Sotiris Ioannidis. 2015. Face/Off: Preventing Privacy Leakage From Photos in Social Networks. In *Proceedings of the 22nd ACM SIGSAC Conference on Computer and Communications Security* (Denver, Colorado, USA) *(CCS '15)*. Association for Computing Machinery, New York, NY, USA, 781–792. doi:10.1145/2810103.2813603

[29] Jindong Jiang, Wafa Skalli, Ali Siadat, and Laurent Gajny. 2022. Effect of Face Blurring on Human Pose Estimation: Ensuring Subject Privacy for Medical and Occupational Health Applications. *Sensors* 22, 23 (2022). doi:10.3390/s22239376

[30] Ji-Eun Joo, Yu Hu, Sujin Kim, Hyunji Kim, Sunyoung Park, Ji-Hoon Kim, Younghyun Kim, and Sung-Min Park. 2022. An Indoor-Monitoring LiDAR Sensor for Patients with Alzheimer Disease Residing in Long-Term Care Facilities. *Sensors* 22, 20 (2022). doi:10.3390/s22207934

[31] Mark Jeffrey Keith, Courtenay Maynes, Paul Benjamin Lowry, and Jeffry Stephen Babb. 2014. Privacy Fatigue: The Effect of Privacy Control Complexity on Consumer Electronic Information Disclosure. *Information Privacy Law eJournal* (2014). https://api.semanticscholar.org/CorpusID:14320957

[32] Injung Kim and Adam J. Lee. 2024. "I know what you did last semester": Understanding Privacy Expectations and Preferences in the Smart Campus. In

Proceedings of the 2024 CHI Conference on Human Factors in Computing Systems (Honolulu, HI, USA) *(CHI '24)*. Association for Computing Machinery, New York, NY, USA, Article 980, 15 pages. doi:10.1145/3613904.3642174

[33] Krishna Kodur and Maria Kyrarini. 2023. Patient–Robot Co-Navigation of Crowded Hospital Environments. *Applied Sciences* 13, 7 (2023). doi:10.3390/app13074576

[34] Lamyanba Laishram, Muhammad Shaheryar, Jong Taek Lee, and Soon Ki Jung. 2025. Toward a Privacy-Preserving Face Recognition System: A Survey of Leakages and Solutions. *ACM Comput. Surv.* 57, 6, Article 147 (Feb. 2025), 38 pages. doi:10.1145/3673224

[35] Adam J. Lee, Jacob T. Biehl, and Conor Curry. 2018. Sensing or Watching? Balancing Utility and Privacy in Sensing Systems via Collection and Enforcement Mechanisms. In *Proceedings of the 23nd ACM on Symposium on Access Control Models and Technologies* (Indianapolis, Indiana, USA) *(SACMAT '18)*. Association for Computing Machinery, New York, NY, USA, 105–116. doi:10.1145/3205977.3205983

[36] Hosub Lee and Alfred Kobsa. 2016. Understanding User Privacy in Internet of Things Environments. In *2016 IEEE 3rd World Forum on Internet of Things (WF-IoT)*. IEEE Press, IEEE Press, Reston, VA, 407–412. doi:10.1109/WF-IoT.2016.7845392

[37] Jun Liu, Amir Shahroudy, Mauricio Perez, Gang Wang, Ling-Yu Duan, and Alex Kot. 2019. NTU RGB+D 120: A Large-Scale Benchmark for 3D Human Activity Understanding. *IEEE Transactions on Pattern Analysis and Machine Intelligence* (05 2019), 1–18. doi:10.1109/TPAMI.2019.2916873

[38] Jun Liu, Amir Shahroudy, Mauricio Perez, Gang Wang, Ling-Yu Duan, and Alex C. Kot. 2020. NTU RGB+D 120: A Large-Scale Benchmark for 3D Human Activity Understanding. *IEEE Transactions on Pattern Analysis and Machine Intelligence* 42, 10 (2020), 2684–2701. doi:10.1109/TPAMI.2019.2916873

[39] Jian-Guang Lou, Hua Cai, and Jiang Li. 2005. A real-time interactive multi-view video system. In *Proceedings of the 13th Annual ACM International Conference on Multimedia* (Hilton, Singapore) *(MULTIMEDIA '05)*. Association for Computing Machinery, New York, NY, USA, 161–170. doi:10.1145/1101149.1101173

[40] Aniello Maiese, Alice Chiara Manetti, Costantino Ciallella, and Vittorio Fineschi. 2022. The Introduction of a New Diagnostic Tool in Forensic Pathology: LiDAR Sensor for 3D Autopsy Documentation. *Biosensors* 12, 2 (2022). doi:10.3390/bios12020132

[41] Karola Marky, Sarah Prange, Florian Krell, Max Mühlhäuser, and Florian Alt. 2020. "You just can't know about everything": Privacy Perceptions of Smart Home Visitors. In *Proceedings of the 19th International Conference on Mobile and Ubiquitous Multimedia* (Essen, Germany) *(MUM '20)*. Association for Computing Machinery, New York, NY, USA, 83–95. doi:10.1145/3428361.3428464

[42] Karola Marky, Alexandra Voit, Alina Stöver, Kai Kunze, Svenja Schröder, and Max Mühlhäuser. 2020. "I don't know how to protect myself": Understanding Privacy Perceptions Resulting from the Presence of Bystanders in Smart Environments. In *Proceedings of the 11th Nordic Conference on Human-Computer Interaction: Shaping Experiences, Shaping Society* (Tallinn, Estonia) *(NordiCHI '20)*. Association for Computing Machinery, New York, NY, USA, Article 4, 11 pages. doi:10.1145/3419249.3420164

[43] Nobuhiro Miyazaki, Kentaro Tsuji, Mingxie Zheng, Moyuri Nakashima, Yuji Matsuda, and Eigo Segawa. 2015. Privacy-conscious human detection using low-resolution video. In *2015 3rd IAPR Asian Conference on Pattern Recognition (ACPR)*. 326–330. doi:10.1109/ACPR.2015.7486519

[44] Md Parvez Mollah, Biplob Debnath, Murugan Sankaradas, Srimat Chakradhar, and Abdullah Mueen. 2025. Roadside Multi-LiDAR Data Fusion for Enhanced Traffic Safety. In *Proceedings of the 31st ACM SIGKDD Conference on Knowledge Discovery and Data Mining V.1* (Toronto ON, Canada) *(KDD '25)*. Association for Computing Machinery, New York, NY, USA, 2469–2479. doi:10.1145/3690624.3709410

[45] Ehsan Nateghinia and Luis F. Miranda-Moreno. 2025. Development of an unsupervised 3D LiDAR-based methodology for automated safety monitoring of railway facilities. *International Journal of Transportation Science and Technology* (2025). doi:10.1016/j.ijtst.2025.01.011

[46] Carman Neustaedter, Saul Greenberg, and Michael Boyle. 2006. Blur filtration fails to preserve privacy for home-based video conferencing. *ACM Trans. Comput.-Hum. Interact.* 13, 1 (March 2006), 1–36. doi:10.1145/1143518.1143519

[47] Elaine Newton, Latanya Sweeney, and Bradley Malin. 2005. Preserving privacy by de-identifying face images. *Knowledge and Data Engineering, IEEE Transactions on* 17 (03 2005), 232– 243. doi:10.1109/TKDE.2005.32

[48] A. T. Nghiem, F. Bremond, M. Thonnat, and V. Valentin. 2007. ETISEO, performance evaluation for video surveillance systems. In *2007 IEEE Conference on Advanced Video and Signal Based Surveillance*. 476–481. doi:10.1109/AVSS.2007.4425357

[49] Wei Niu, Jiao Long, Dan Han, and Yuan-Fang Wang. 2004. Human activity detection and recognition for video surveillance. In *2004 IEEE International Conference on Multimedia and Expo (ICME) (IEEE Cat. No.04TH8763)*, Vol. 1. 719–722 Vol.1. doi:10.1109/ICME.2004.1394293

[50] Matija Novak, Ruben Picek, and Darko Andročec. 2025. Overview of Smart Home Security with Raspberry. In *Proceedings of the 2024 7th International Conference on Information Science and Systems (ICISS '24)*. Association for Computing Machinery, New York, NY, USA, 163–169. doi:10.1145/3700706.3700733

[51] Masakazu Ohno, Riki Ukyo, Tatsuya Amano, Hamada Rizk, and Hirozumi Yamaguchi. 2023. Privacy-preserving Pedestrian Tracking using Distributed 3D LiDARs. In *2023 IEEE International Conference on Pervasive Computing and Communications (PerCom)*. 43–52. doi:10.1109/PERCOM56429.2023.10099061

[52] José Ramón Padilla-López, Alexandros Andre Chaaraoui, and Francisco Flórez-Revuelta. 2015. Visual privacy protection methods: A survey. *Expert Systems with Applications* 42, 9 (2015), 4177–4195. doi:10.1016/j.eswa.2015.01.041

[53] Jessica Percy Campbell, Jacob Buchan, Charlene H Chu, Andria Bianchi, Jesse Hoey, and Shehroz S Khan. 2024. User Perception of Smart Home Surveillance Among Adults Aged 50 Years and Older: Scoping Review. *JMIR Mhealth Uhealth* 12 (9 Feb 2024), e48526. doi:10.2196/48526

[54] Francesco Pistolesi, Michele Baldassini, and Beatrice Lazzerini. 2024. A LiDAR-Based Recommendation System to Improve Sitting Posture Respecting Privacy. *IEEE Journal of Biomedical and Health Informatics* 28, 9 (2024), 5303–5311. doi:10.1109/JBHI.2024.3404233

[55] Ramesh Chandra Poonia and Halah A. Al-Alshaikh. 2024. Ensemble approach of transfer learning and vision transformer leveraging explainable AI for disease diagnosis: An advancement towards smart healthcare 5.0. *Computers in Biology and Medicine* 179 (2024), 108874. doi:10.1016/j.compbiomed.2024.108874

[56] Omar Rinchi, Ahmad Alsharoa, and Denise A. Baker. 2024. Remote Breathing Monitoring Using LiDAR Technology. In *2024 46th Annual International Conference of the IEEE Engineering in Medicine and Biology Society (EMBC)*. 1–6. doi:10.1109/EMBC53108.2024.10782856

[57] Omar Rinchi, Nathanael Nisbett, and Ahmad Alsharoa. 2023. Patients Arms Segmentation and Gesture Identification Using Standalone 3-D LiDAR Sensors. *IEEE Sensors Letters* 7, 9 (2023), 1–4. doi:10.1109/LSENS.2023.3303081

[58] Bruno Rodrigues, Lukas Müller, Eder J. Scheid, Muriel F. Franco, Christian Killer, and Burkhard Stiller. 2021. LaFlector: a Privacy-preserving LiDAR-based Approach for Accurate Indoor Tracking. In *2021 IEEE 46th Conference on Local Computer Networks (LCN)*. 367–370. doi:10.1109/LCN52139.2021.9524945

[59] Michael S. Ryoo, Kiyoon Kim, and Hyun Jong Yang. 2018. Extreme low resolution activity recognition with multi-siamese embedding learning. In *Proceedings of the Thirty-Second AAAI Conference on Artificial Intelligence and Thirtieth Innovative Applications of Artificial Intelligence Conference and Eighth AAAI Symposium on Educational Advances in Artificial Intelligence* (New Orleans, Louisiana, USA) *(AAAI'18/IAAI'18/EAAI'18)*. AAAI Press, Article 896, 8 pages.

[60] Michael S. Ryoo, Brandon Rothrock, Charles Fleming, and Hyun Jong Yang. 2017. Privacy-preserving human activity recognition from extreme low resolution. In *Proceedings of the Thirty-First AAAI Conference on Artificial Intelligence* (San Francisco, California, USA) *(AAAI'17)*. AAAI Press, 4255–4262.

[61] Aaqib Saeed, Tanir Ozcelebi, and Johan Lukkien. 2019. Multi-task Self-Supervised Learning for Human Activity Detection. *Proc. ACM Interact. Mob. Wearable Ubiquitous Technol.* 3, 2, Article 61 (June 2019), 30 pages. doi:10.1145/3328932

[62] Mukesh Saini, Pradeep K. Atrey, Sharad Mehrotra, and Mohan Kankanhalli. 2014. W3-privacy: understanding what, when, and where inference channels in multi-camera surveillance video. *Multimedia Tools Appl.* 68, 1 (Jan. 2014), 135–158. doi:10.1007/s11042-012-1207-9

[63] Md. Haidar Sharif. 2021. Laser-Based Algorithms Meeting Privacy in Surveillance: A Survey. *IEEE Access* 9 (2021), 92394–92419. doi:10.1109/ACCESS.2021.3092687

[64] Xia Su, Han Zhang, Kaiming Cheng, Jaewook Lee, Qiaochu Liu, Wyatt Olson, and Jon E. Froehlich. 2024. RASSAR: Room Accessibility and Safety Scanning in Augmented Reality. In *Proceedings of the 2024 CHI Conference on Human Factors in Computing Systems* (Honolulu, HI, USA) *(CHI '24)*. Association for Computing Machinery, New York, NY, USA, Article 968, 17 pages. doi:10.1145/3613904.3642140

[65] Jaeyong Sung, Colin Ponce, Bart Selman, and Ashutosh Saxena. 2012. Unstructured human activity detection from RGBD images. In *2012 IEEE International Conference on Robotics and Automation*. 842–849. doi:10.1109/ICRA.2012.6224591

[66] Karl van der Schyff, Greg Foster, Karen Renaud, and Stephen Flowerday. 2023. Online Privacy Fatigue: A Scoping Review and Research Agenda. *Future Internet* 15, 5 (2023). doi:10.3390/fi15050164

[67] Richmond Y. Wong, Jason Caleb Valdez, Ashten Alexander, Ariel Chiang, Olivia Quesada, and James Pierce. 2023. Broadening Privacy and Surveillance: Eliciting Interconnected Values with a Scenarios Workbook on Smart Home Cameras. In *Proceedings of the 2023 ACM Designing Interactive Systems Conference* (Pittsburgh, PA, USA) *(DIS '23)*. Association for Computing Machinery, New York, NY, USA, 1093–1113. doi:10.1145/3563657.3596012

[68] Xuan Wu, Haiyang Zhang, Chunxiu Kong, Yuanze Wang, Yezhao Ju, and Changming Zhao. 2024. LiDAR-Based 3-D Human Pose Estimation and Action Recognition for Medical Scenes. *IEEE Sensors Journal* 24, 9 (2024), 15531–15539. doi:10.1109/JSEN.2024.3373192

[69] Mingze Xu, Aidean Sharghi, Xin Chen, and David J. Crandall. 2018. Fully-Coupled Two-Stream Spatiotemporal Networks for Extremely Low Resolution Action Recognition. In *2018 IEEE Winter Conference on Applications of Computer Vision (WACV)*. 1607–1615. doi:10.1109/WACV.2018.00178

[70] Mingze Xu, Aidean Sharghi, Xin Chen, and David J. Crandall. 2018. Fully-Coupled Two-Stream Spatiotemporal Networks for Extremely Low Resolution Action

Recognition. In *2018 IEEE Winter Conference on Applications of Computer Vision (WACV)*. 1607–1615. doi:10.1109/WACV.2018.00178

[71] Ronghua Xu, Seyed Yahya Nikouei, Yu Chen, Aleksey Polunchenko, Sejun Song, Chengbin Deng, and Timothy R. Faughnan. 2018. Real-Time Human Objects Tracking for Smart Surveillance at the Edge. In *2018 IEEE International Conference on Communications (ICC)*. 1–6. doi:10.1109/ICC.2018.8422970

[72] Yaxing Yao, Justin Reed Basdeo, Oriana Rosata Mcdonough, and Yang Wang. 2019. Privacy Perceptions and Designs of Bystanders in Smart Homes. *Proc. ACM Hum.-Comput. Interact.* 3, CSCW, Article 59 (Nov. 2019), 24 pages. doi:10.1145/3359161

[73] Shikun Zhang, Yuanyuan Feng, Lujo Bauer, Lorrie Cranor, Anupam Das, and Norman Sadeh. 2021. "Did you know this camera tracks your mood?": Understanding Privacy Expectations and Preferences in the Age of Video Analytics. *Proceedings on Privacy Enhancing Technologies* 2021 (04 2021), 282–304. doi:10.2478/popets-2021-0028

[74] Serena Zheng, Noah Apthorpe, Marshini Chetty, and Nick Feamster. 2018. User Perceptions of Smart Home IoT Privacy. *Proc. ACM Hum.-Comput. Interact.* 2, CSCW, Article 200 (Nov. 2018), 20 pages. doi:10.1145/3274469

[75] Yin Zhou and Oncel Tuzel. 2018. VoxelNet: End-to-End Learning for Point Cloud Based 3D Object Detection. In *2018 IEEE/CVF Conference on Computer Vision and Pattern Recognition*. 4490–4499. doi:10.1109/CVPR.2018.00472

11 APPENDIX

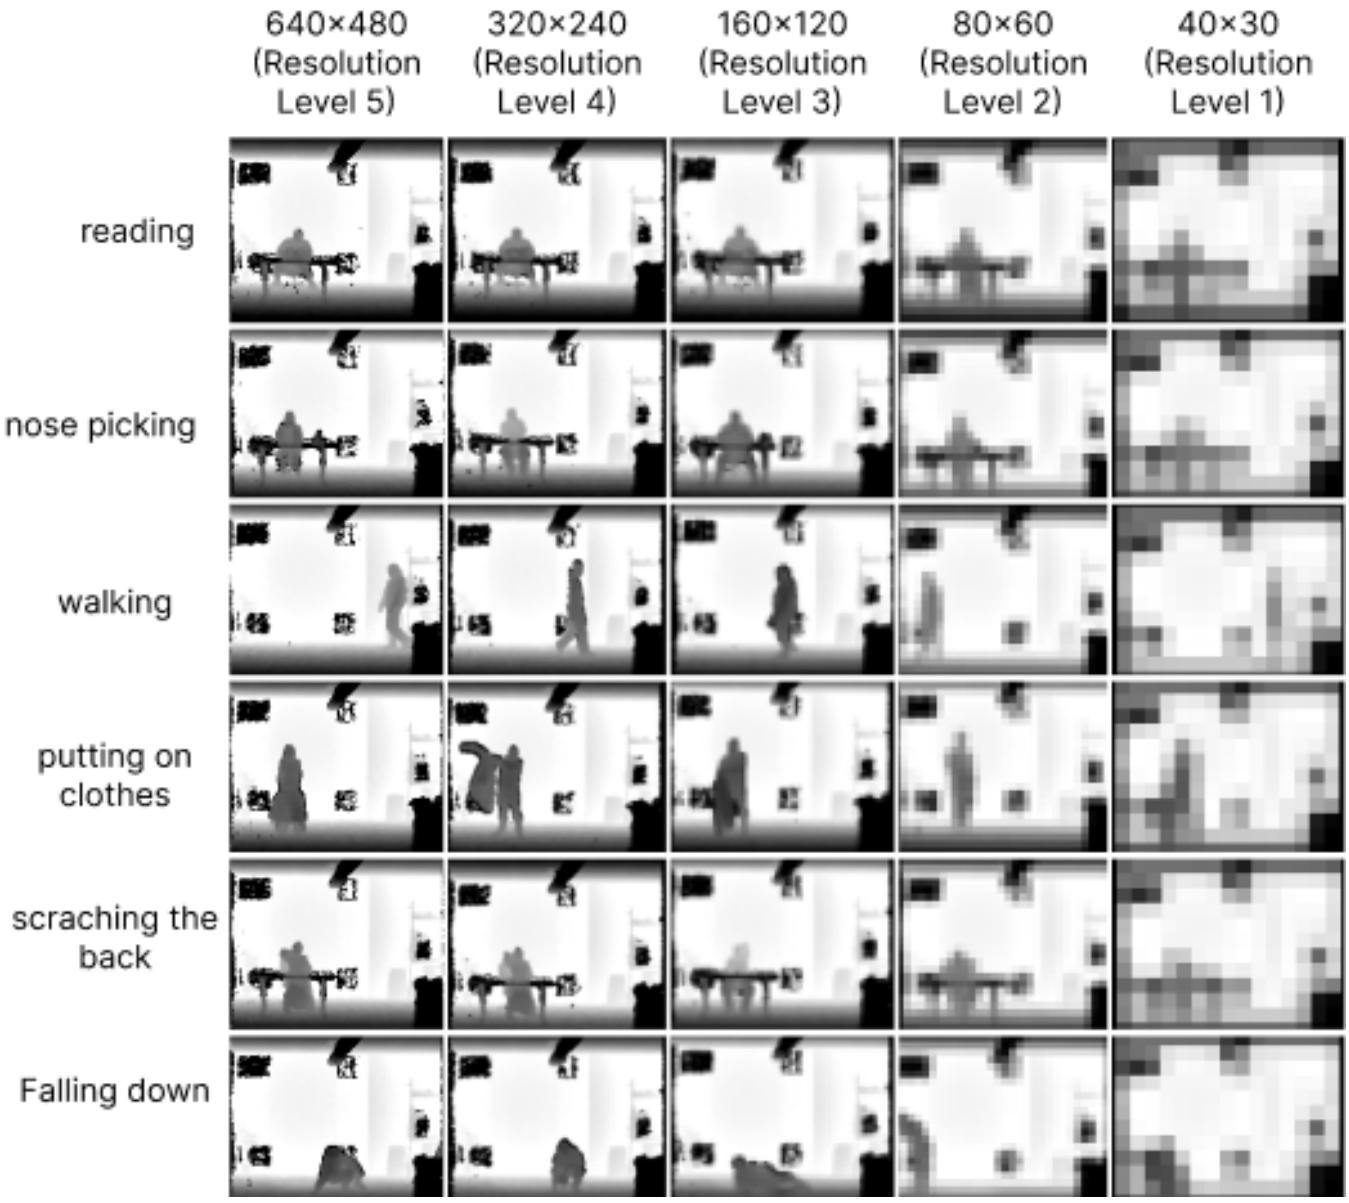


Figure 18: Sample images from action data collection, which list the images at five different resolution levels for all actions.

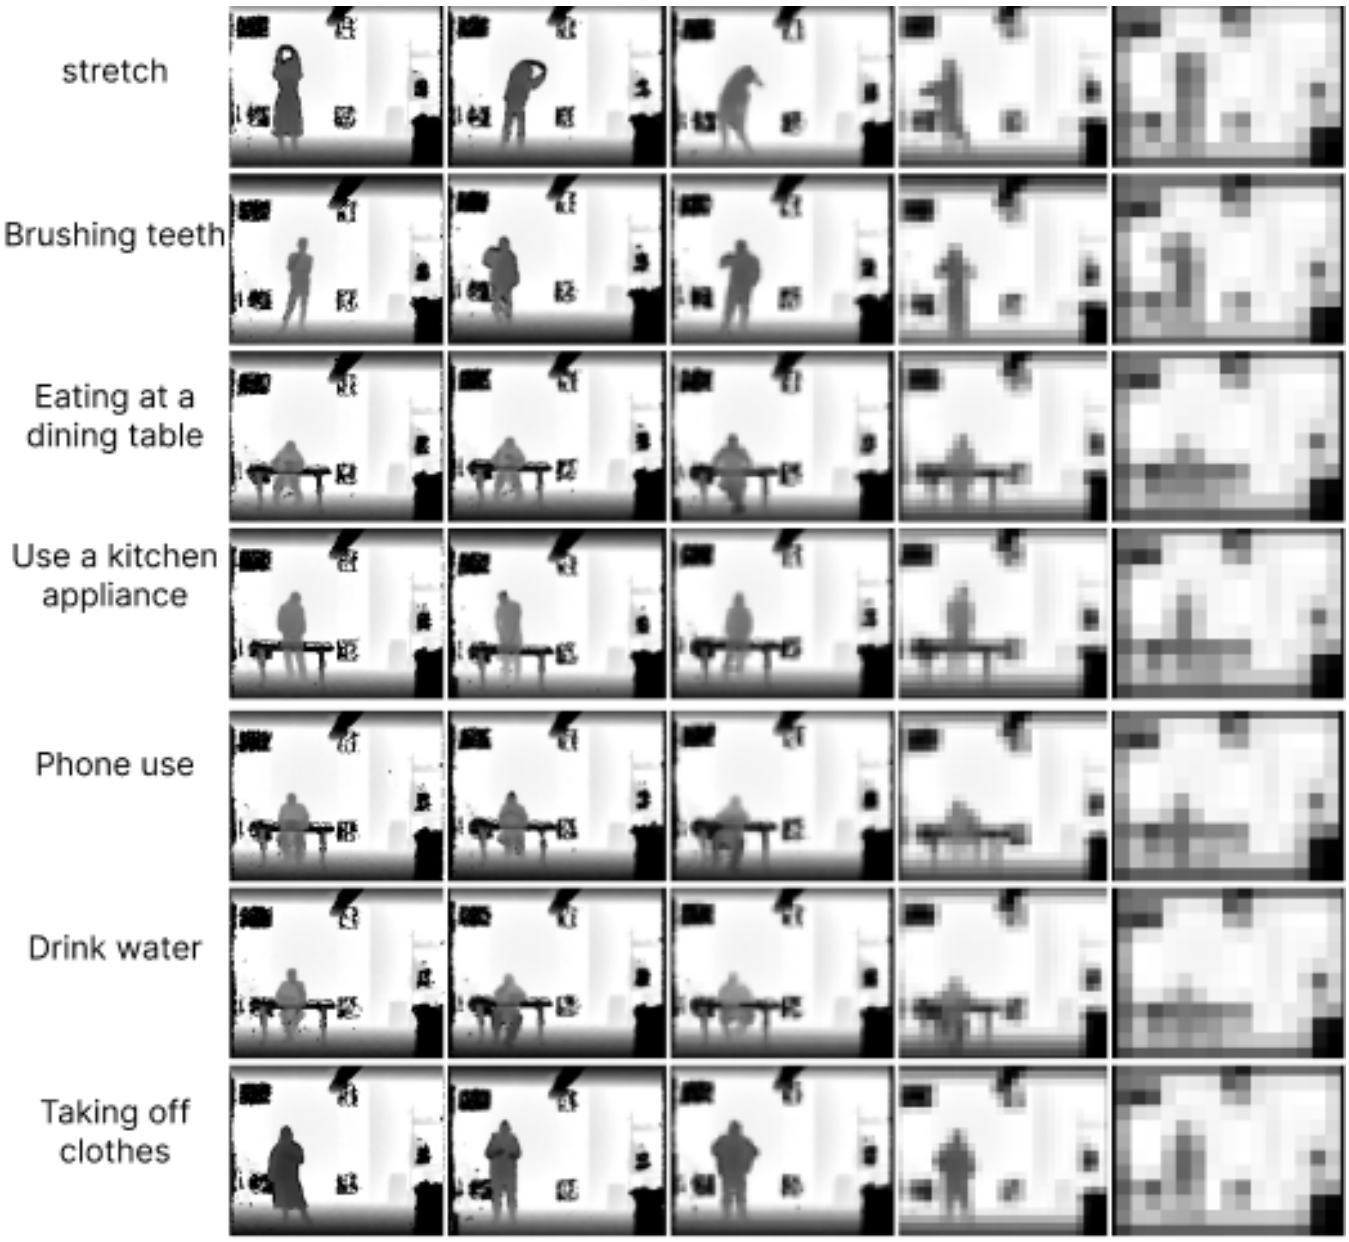


Figure 19: Sample images from action data collection, which list the images at five different resolution levels for all actions.

Class	Precision	Recall	F1	Resolution Level
Brushing teeth	0.99	0.97	0.98	640×480
	0.95	0.90	0.92	320×240
	0.91	0.83	0.87	160×120
	0.92	0.56	0.70	80×60
	0.64	0.67	0.66	40×30
Drink water	1.00	0.88	0.93	640×480
	1.00	0.80	0.89	320×240
	0.96	0.88	0.92	160×120
	0.81	0.84	0.83	80×60
	0.92	0.52	0.67	40×30
Eating at a dining table	1.00	0.93	0.97	640×480
	0.96	0.85	0.90	320×240
	0.92	0.81	0.86	160×120
	0.61	0.37	0.46	80×60
	0.25	0.59	0.36	40×30
Falling down	1.00	1.00	1.00	640×480
	1.00	1.00	1.00	320×240
	0.96	0.93	0.94	160×120
	1.00	0.93	0.96	80×60
	0.98	0.86	0.91	40×30
Reading at a desk	0.93	0.94	0.94	640×480
	0.97	0.85	0.91	320×240
	0.80	0.78	0.79	160×120
	0.76	0.52	0.62	80×60
	0.43	0.67	0.52	40×30
Use a kitchen appliance	0.93	0.92	0.93	640×480
	0.90	0.83	0.87	320×240
	0.88	0.66	0.75	160×120
	0.64	0.67	0.65	80×60
	0.47	0.05	0.09	40×30

Table 2: Evaluation results for the first 6 actions

Class	Precision	Recall	F1	Resolution Level
Walking	1.00	1.00	1.00	640×480
	1.00	1.00	1.00	320×240
	1.00	1.00	1.00	160×120
	1.00	1.00	1.00	80×60
	0.99	0.88	0.93	40×30
Taking off clothes	0.88	0.98	0.92	640×480
	0.71	0.96	0.82	320×240
	0.79	0.93	0.85	160×120
	0.57	0.92	0.71	80×60
	0.50	0.89	0.64	40×30
Putting on clothes	0.95	1.00	0.97	640×480
	0.86	0.89	0.87	320×240
	0.71	0.96	0.82	160×120
	0.37	0.51	0.43	80×60
	0.28	0.14	0.18	40×30
Nose picking	1.00	1.00	1.00	640×480
	1.00	0.99	0.99	320×240
	0.96	1.00	0.98	160×120
	0.80	0.91	0.86	80×60
	0.48	0.89	0.63	40×30
Phone use	0.94	0.99	0.97	640×480
	0.89	0.99	0.94	320×240
	0.88	0.97	0.92	160×120
	0.88	0.45	0.59	80×60
	0.64	0.46	0.54	40×30
Stretch	0.97	0.95	0.96	640×480
	0.99	0.97	0.98	320×240
	1.00	0.88	0.94	160×120
	0.85	0.77	0.81	80×60
	0.46	0.85	0.60	40×30
Scratching the back	0.95	0.99	0.97	640×480
	0.85	0.98	0.91	320×240
	0.87	0.89	0.88	160×120
	0.52	0.87	0.66	80×60
	0.61	0.27	0.38	40×30

Table 3: Evaluation results for the remaining 7 actions