

Lightweight Range–Angle Imaging Based Algorithm for Quasi-Static Human Detection on Low-Cost FMCW Radar

Huy Trinh
University of Waterloo
h3trinh@uwaterloo.ca

George Shaker
University of Waterloo
gshaker@uwaterloo.ca

Abstract

Quasi-static human activities such as lying, standing or sitting produce very low Doppler shifts and highly spread radar signatures, making them difficult to detect with conventional constant-false-alarm rate (CFAR) detectors tuned for point targets. Moreover, privacy concerns and low lighting conditions limit the use of cameras in long-term care (LTC) facilities. This paper proposes a lightweight, non-visual image-based method for robust quasi-static human presence detection using a low-cost 60 GHz FMCW radar. On a dataset covering five semi-static activities, the proposed method improves average detection accuracy from 68.3% for Cell-Averaging CFAR (CA-CFAR) and 78.8% for Order-Statistics CFAR (OS-CFAR) to 93.24% for Subject 1, from 51.3%, 68.3% to 92.3% for Subject 2, and 57.72%, 69.94% to 94.82% for Subject 3, respectively. Finally, we benchmarked all three detectors across all activities on a Raspberry Pi 4B using a shared Range-Angle (RA) preprocessing pipeline. The proposed algorithm obtains an average 8.2 ms per frame, resulting in over 120 frames per second (FPS) and a 74× speed-up over OS-CFAR. These results demonstrate that simple image-based processing can provide robust and deployable quasi-static human sensing in cluttered indoor environments.

1. Introduction

Semi-static human activities such as lying on a sofa, sitting, or standing are common in long-term-care (LTC) facilities, yet they remain difficult to monitor with conventional radar processing [9, 15]. Recent remote-sensing work has focused either on range–Doppler or range–angle feature maps and deep neural networks trained on high-resolution radar data. Christian et al. [4] evaluate several image transformations of range–Doppler maps and study their impact on lightweight convolutional models for person localization, while Stephan et al. [16] propose a deep-learning architecture to detect people in range–Doppler images. These meth-

ods, however, do not explicitly target quasi-static occupants with very low radial velocity. On the other hand, several studies have explored treating radar feature maps as images and applying image processing and morphological operations on range–angle maps to estimate angle-of-arrival and improve visual representation [5, 12], and other authors study YOLO-based detectors on range–angle heatmaps to detect multiple targets under various scenarios [6, 9]. However, these approaches either rely on heavy deep networks or assume high-end MIMO arrays [8], which limits their suitability for low-power embedded deployment in LTC rooms. To the best of our knowledge, no prior work has proposed a lightweight image-based method specifically tailored to quasi-static human presence detection using low-cost FMCW radar with only a few receive channels. In this paper we bridge radar signal processing and image processing by: (i) treating the Capon-based range–azimuth (RA) map as a 2-D image; (ii) systematically evaluating 2-D CA-CFAR and OS-CFAR on these RA images under our real-world high-clutter condition dataset; and (iii) introducing a simple percentile-gated detector on range–angle to improve semi-static detection accuracy and lightweight for edge deployment.

2. Methodologies

2.1. Data Collecting and Preprocessing

For this study, we collected a dedicated dataset of semi-static activities in a LTC using an off-the-shelf 60 GHz mmWave FMCW radar (Infineon XENSIV™ BGT60TR13C) with one transmitter and three receivers [7]. The radar was wall-mounted at approximately 2.5 m height with partially overlapping fields of view and a slight downward tilt to cover both sofa and floor areas. The room contains typical living-room furniture and occlusions (sofa, coffee table, TV, windows, and surrounding walls), providing realistic multipath and clutter. For each activity (standing, sitting on the sofa, sitting on floor, lying on the sofa, lying on the floor), three subjects were recorded for two minutes each at multiple aspect angles and ranges at 10 Hz rate.

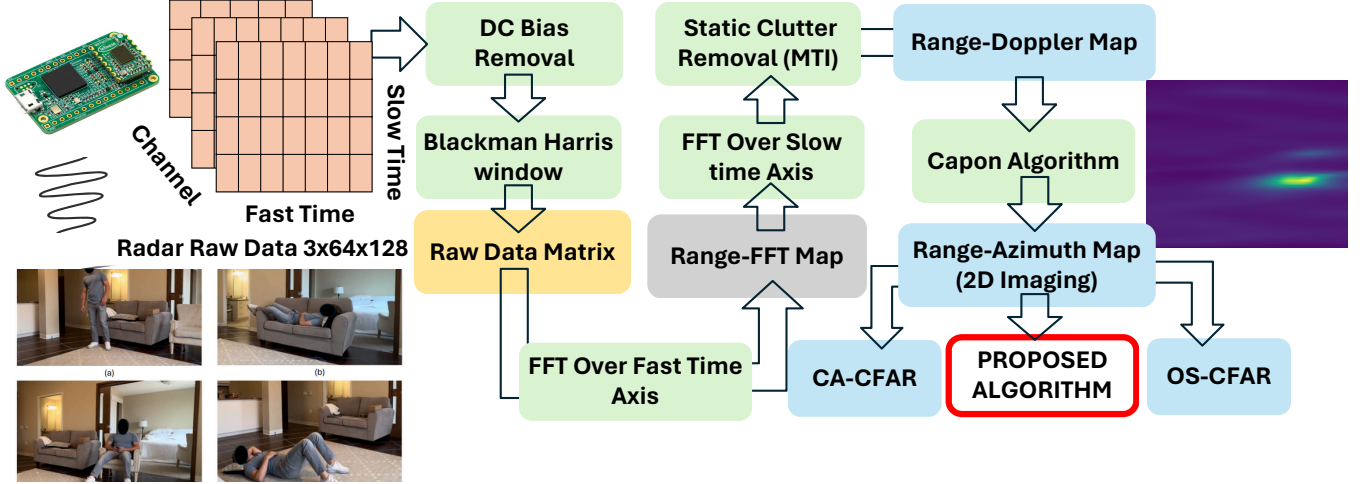


Figure 1. End-to-end radar’s imaging processing pipeline. Raw complex IQ frames from the 60 GHz FMCW radar are converted to range–Doppler maps via range and Doppler FFTs with DC removal, windowing, and MTI clutter suppression. Capon beamforming over two receive channels produces a normalized range–azimuth (RA) intensity map, which is then processed by three detectors: 2-D CA–CFAR, 2-D OS–CFAR, and the proposed percentile-gated lump detector.

Each radar frame provides a complex-valued data cube of size $3 \times N_{\text{chirp}} \times N_{\text{sample}} = 3 \times 128 \times 64$. Our preprocessing pipeline follows Infineon baseline [3] and summarized in Fig. 1. To suppress static clutter we apply an exponential moving-average high-pass filter in the RD domain, similar to a moving-target-indicator (MTI) stage [17, 19]. Finally, range-angle feature maps are formed using a Capon beamformer across the two spatially separated receive channels (RX0 and RX2). For each range bin, we estimate a 2×2 spatial covariance matrix from Doppler bins and compute the Capon spatial spectrum over azimuth [1, 2, 18]. Stacking the resulting power values over range and azimuth yields the RA map $X \in \mathbb{R}^{H \times W}$ (with $H = 64$ ranges and $W = 256$ azimuth bins). Each frame is normalized by its maximum value so that X lies in $[0, 1]$ as visualized in Fig. 2. These RA images are common inputs to all detection methods used in the remainder of this work.

2.2. Baseline 2-D CA–CFAR and OS–CFAR

Given a range–azimuth image X , conventional 2-D CFAR operates on each cell-under-test (CUT) by comparing its power $|X(\text{CUT})|^2$ to an adaptive threshold derived from neighbouring “reference” cells [10, 11]. The detection decision can be written as

$$I_{\text{det}}(\text{CUT}) = \begin{cases} 1, & |X(\text{CUT})|^2 > \mu \hat{\sigma}^2, \\ 0, & \text{otherwise,} \end{cases} \quad (1)$$

where $\hat{\sigma}^2$ is an estimate of the local noise-plus-clutter power and μ is a scale factor chosen to achieve a desired probability of false alarm P_{fa} . In **cell averaging CFAR (CA–CFAR)**, $\hat{\sigma}^2$ is obtained by averaging the power in a

Range-Azimuth Surface Response

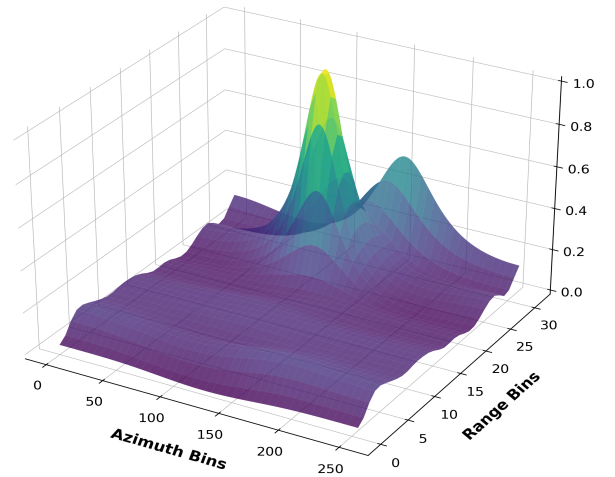


Figure 2. Capon-based range–azimuth surface response (1 frame) for a subject lying on the sofa. The subject produces a compact high-intensity peak in range–azimuth space on top of a broad background clutter ridge from walls and furniture.

sliding $K \times K$ window of reference cells surrounding the CUT, excluding a guard region directly around the CUT. For a given P_{fa} and N reference cells, the scale factor is [19]

$$\mu = N(P_{\text{fa}}^{-1/N} - 1), \quad (2)$$

and the threshold $\mu \hat{\sigma}^2$ is updated at each pixel as the window moves across the RA image. **Order statistics CFAR (OS–CFAR)** [14] replaces the mean of the reference win-




Figure 3. Percentile ablation on a representative range–azimuth frame (*lying on sofa*). From left to right: normalized RA image and detection masks obtained by keeping only the top 90%, 95%, and 99% of pixel intensities.

dow by an order statistic that is more robust to outliers: the reference powers are sorted and the k -th largest value is used as $\hat{\sigma}^2$ [19].

2.3. Proposed Range–Angle Subject Detector

Instead of applying CFAR directly on the RA image, we treat X as a 2-D grayscale image and search for a single connected lump that corresponds to the human reflector. Our detector focuses on the top p -percentile pixel intensities and enforces spatial coherence.

Algorithm 1 Range–Angle Lump Detector

```

1:  $\tau \leftarrow \text{PERCENTILE}(X, p)$  {gate top- $p$  intensity}
2:  $M \leftarrow 0_{H \times W}$ 
3: for all pixels  $(i, j)$  do
4:   if  $X[i, j] \geq \tau$  then
5:      $M[i, j] \leftarrow 1$  {keep bright responses}
6:   end if
7: end for
8:  $M \leftarrow \text{Close}(M)$  {connect nearby pixels (small structuring element)}
9:  $M \leftarrow \text{RemoveSmall}(M, A_{\min})$  {discard tiny blobs}
10:  $\mathcal{C} \leftarrow \text{ConnectedComponents}(M)$ 
11: if  $|\mathcal{C}| = 0$  then
12:   return  $M = 0_{H \times W}, B = \emptyset$ 
13: end if
14:  $c^* \leftarrow \arg \max_{c \in \mathcal{C}} \text{Area}(c)$  {dominant lump}
15:  $M \leftarrow 0_{H \times W}$ ;
16: for all pixels  $(i, j)$  in  $c^*$  do
17:    $M[i, j] \leftarrow 1$ 
18: end for
19:  $B \leftarrow \text{BBOX}(c^*)$ 
20:
21: return  $M, B$ 

```

Algorithm 1 summarizes our proposed procedure. Given a normalized RA image X , we first compute a global percentile threshold τ and form a binary mask M that keeps only pixels above τ . This step removes most of the floor and wall clutter while preserving the main lobe of the human response. Next, we apply a small 2×2 structuring element for binary closing to connect nearby bright pixels

along range and azimuth, and we remove any blobs whose area is below A_{\min} . A two–pass connected–components algorithm labels the remaining blobs, and we keep only the largest component as the final detection lump. All operations in Algorithm 1 are linear in the number of pixels and involve only simple comparisons and logic, resulting in $\mathcal{O}(HW)$ complexity per frame. In contrast, 2-D CA-CFAR and OS-CFAR require sliding windows over each pixel with relatively large kernels, resulting in $\mathcal{O}(HWK^2)$ complexity and substantially higher runtime. Figure 3 illustrates that different percentile intensity p affects the detection mask on a representative RA frame for the lying on sofa activity of subject 1. At the 90th and 95th percentiles, the mask still includes a wide horizontal band along the dominant clutter ridge, causing the detected lump to spread across the entire sofa region. In contrast, the 99th percentile isolates a smaller, well-localized lump that aligns with the annotated human activity detection.

3. Results and Discussion

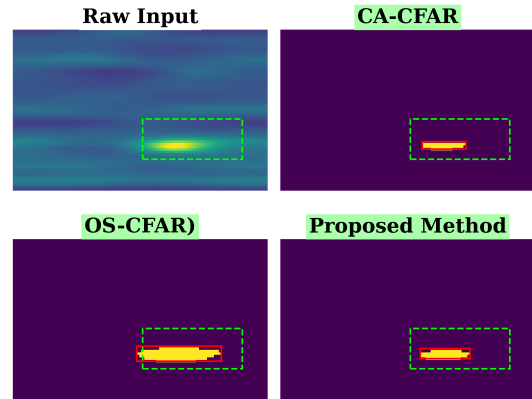


Figure 4. Qualitative comparison of all three detectors on frame 21 for subject 1 lying on the floor, which all of them produce detections within the ground-truth bounding box.

Figures 4 and 5 compare the detection results on representative frames for Subject 1 lay on sofa. In these visualizations, the subject’s actual position is annotated by a green dashed bounding box. We further encode the decision out-

Table 1. Detection accuracy (%) per subject. Best values for each metric per subject are shown in **bold**.

Activity	Subject 1			Subject 2			Subject 3		
	CA	OS	Ours	CA	OS	Ours	CA	OS	Ours
Standing	79.0	89.9	93.5	45.0	65.7	91.7	80.0	91.2	97.3
Sit on Sofa	75.7	88.2	97.0	31.4	51.6	86.9	57.0	72.2	92.7
Sit on Floor	83.3	91.5	97.2	54.6	68.4	94.1	74.7	82.9	97.4
Lay on Floor	70.6	80.5	91.9	63.9	77.0	94.7	78.4	86.4	93.7
Lay on Sofa	33.0	43.9	86.6	61.8	79.0	94.1	48.5	67.0	93.0
<i>Average</i>	68.3	78.8	93.24	51.3	68.3	92.3	57.72	69.94	94.82

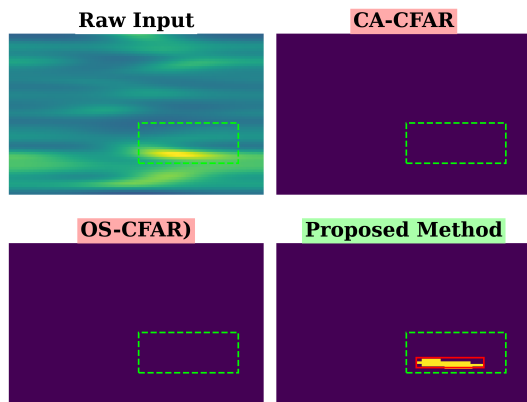


Figure 5. Qualitative comparison of all three detectors on frame 106 for subject 1 lying on the floor. only our proposed method can detect within the ground-truth bounding box.

come in the method headers: we count a frame as a hit if one detected blob overlaps the annotated ground-truth bounding box. A green highlight indicates a successful detection, while a red highlight signifies a miss. While all methods successfully localize the subject within the ground-truth region during the high signal-to-noise ratio (SNR) scenario (Frame 21, Fig. 4), the proposed algorithm demonstrates superior stability under severe noise. As shown in Frame 106 (Fig. 5), strong background clutter causes both CA-CFAR and OS-CFAR to fail, whereas our proposed method succeeds in detecting the subject within the ground-truth bounds. Table 1 reports per-activity detection accuracy for three subjects. On average, the proposed method improves detection accuracy to 93.24% (Subject 1), 92.3% (Subject 2), and 94.82% (Subject 3), significantly outperforming CA-CFAR (68.3%/51.3%/57.72%) and OS-CFAR (78.8%/68.3%/69.94%). and 43.9% (OS) to 86.6% for Subject 1, and 48.5% (CA), 67.0% (OS) to 93% for Subject 3. For Subject 2, the increase is from 31.4% and 51.6% to 86.9% for Subject 2 in Sit on Sofa. Notably, sofa-based activities remain the most difficult, representing the lowest detection performance for Subject 1 (lying) and Subject 2 (sitting), respectively. Finally, to assess real-world feasi-




Figure 6. Raspberry Pi 4B Board with BGT60TR13C radar.

bility and performance, we benchmark all three detectors across all activities using a streaming implementation on a Raspberry Pi 4B [13] as Fig. 6, and the results are shown in Table 2. Range-azimuth feature maps are precomputed using the shared pipeline and then fed to each detector. Compared to CA-CFAR (91.2 ms) and OS-CFAR (601.1 ms), our proposed method runs on average of just 8.2 ms per frame, achieving a throughput of over 122 FPS. This represents a 74 \times speed-up over OS-CFAR and ensures the system easily meets the 10 Hz radar frame-rate constraint without deadline misses. These results confirm that the proposed image-based approach delivers both robust detection accuracy and real-time computational efficiency, suitable for the low-latency requirements of edge deployment.

Table 2. Average runtime analysis on Raspberry Pi 4 for 3 methods. Best values for each metric are shown in **bold**.

Method	Latency (ms)	FPS	RAM (MB)	Speedup (vs. OS)
CA-CFAR	91.2	11.0	299.4	6.6 \times
OS-CFAR	601.1	1.7	299.5	1.0 \times
Ours	8.2	122.4	299.5	73.6\times

4. Conclusion and Future Work

This paper presented a non-visual sensing solution for quasi-static human presence in cluttered indoor environments using a single low-cost 60 GHz FMCW radar. By treating Capon-based range-azimuth maps as 2-D images and applying a lightweight percentile-gated detector, we achieved robust detection where conventional CFAR methods struggle, particularly in high-clutter scenarios. The method relies on basic image-processing primitives, ensuring it is computationally efficient for real-time execution on embedded platforms. Future work will investigate deeper optimization of the processing pipeline to further enhance detection reliability and stability. A key focus will be on the scalability of this edge-based solution. Given the low computational footprint and hardware cost, we aim to validate the system's deployment in large-scale facilities—scaling to scenarios in hospitals and long-term care centers, where privacy-preserving, decentralized, and cost-effective monitoring are critical.

References

- [1] Hajar Abedi, Shenghang Luo, and George Shaker. On the use of low-cost radars and machine learning for in-vehicle passenger monitoring. *2020 IEEE 20th Topical Meeting on Silicon Monolithic Integrated Circuits in RF Systems (SiRF)*, pages 63–65, 2020. 2
- [2] Hajar Abedi, Ahmad Ansariyan, and George Shaker. Contactless in-bed detection using a low-cost low-resolution radar. In *2024 IEEE SENSORS*, pages 1–4, 2024. 2
- [3] Infineon Technologies AG. Ceiling-mounted occupancy detection using xensiv™ demo bgt60tr13c 60 ghz radar. Application Note AN141319 AN141319, Infineon, 2024. Available: <https://www.infineon.com/assets/row/public/documents/24/42/infineon-an141319-ceiling-mounted-occupancy-detection-using-xensiv-demo-bgt60tr13c-60-ghz-radar-applicationnotes-en.pdf>. 2
- [4] Christian Gianoglio, Ammar Mohanna, Ali Rizik, Laurence Moroney, and Maurizio Valle. On edge human action recognition using radar-based sensing and deep learning. *IEEE Transactions on Industrial Informatics*, 20(3):4160–4172, 2024. 1
- [5] Khushi Gupta, Soumya Joshi, M. B. Srinivas, Srinivas Boppu, M. Sabarimalai Manikandan, and Linga Reddy Cenkeramaddi. Localization of multi-class on-road and aerial targets using mmwave fmcw radar. *Electronics*, 10(23), 2021. 1
- [6] Siddharth Gupta, Prabhat Kumar Rai, Abhinav Kumar, Phaneendra K. Yalavarthy, and Linga Reddy Cenkeramaddi. Target classification by mmwave fmcw radars using machine learning on range-angle images. *IEEE Sensors Journal*, 21(18):19993–20001, 2021. 1
- [7] Infineon Technologies AG. Bgt60tr13c 60 ghz radar sensor. *Infineon Product Page*, 2024. [Online]. Available: <https://www.infineon.com/cms/en/product/sensor/radar-sensors/radar-sensors-for-iot/60ghz-radar/bgt60tr13c/> (Accessed: April 15, 2025). 1
- [8] Kevin Kaiser, Thomas Stadelmayr, Jens Felgentreff, Robert Weigel, Fabian Lurz, and Avik Santra. Improved indoor semi-static human target detection and localization using fmcw radar. In *2023 20th European Radar Conference (EuRAD)*, pages 298–301, 2023. 1
- [9] Hongchun Li, Lili Xie, Qian Zhao, Jun Tian, and Takeshi Konno. Static human localization using fmcw mimo radar. In *2023 IEEE Wireless Communications and Networking Conference (WCNC)*, pages 1–6, 2023. 1
- [10] Robert Miller. Fundamentals of radar signal processing (richards, m.a.; 2005) [book review]. *IEEE Signal Processing Magazine*, 26:100–101, 2009. 2
- [11] Minh Q. Nguyen, Reinhard Feger, Thomas Wagner, and Andreas Stelzer. Analysis of 2d ca-cfar for ddma fmcw mimo radar. In *2023 20th European Radar Conference (EuRAD)*, pages 423–426, 2023. 2
- [12] Aarav Pandya, Ajit Jha, and Linga Reddy Cenkeramaddi. A velocity estimation technique for a monocular camera using mmwave fmcw radars. *Electronics*, 10(19), 2021. 1
- [13] Raspberry Pi Ltd. Raspberry pi 4 model b. *Raspberry Pi Product Page*, 2019. [Online]. Available: <https://www.raspberrypi.com/products/raspberry-pi-4-model-b/> (Accessed: April 15, 2025). 4
- [14] Hermann Rohling. Radar cfar thresholding in clutter and multiple target situations. *IEEE Transactions on Aerospace and Electronic Systems*, AES-19(4):608–621, 1983. 2
- [15] Avik Santra and Souvik Hazra. *Deep Learning Applications of Short-Range Radars*. 2020. 1
- [16] Michael Stephan, Avik Santra, and Georg Fischer. *Human Target Detection and Localization with Radars Using Deep Learning*, pages 173–197. Springer Singapore, Singapore, 2021. 1
- [17] Huy Trinh, Elliot Creager, and George Shaker. Doppler-domain respiratory amplification for semi-static human occupancy detection using low-resolution simo fmcw radar, 2026. 2
- [18] Joshua Visser, Hajar Abedi, Sachin Karmani, and George Shaker. Ai based activity monitoring in washrooms using low-resolution radar. In *2025 IEEE MTT-S Latin America Microwave Conference (LAMC)*, pages 78–81, 2025. 2
- [19] Christoph Will, Prachi Vaishnav, Abhiram Chakraborty, and Avik Santra. Human target detection, tracking, and classification using 24-ghz fmcw radar. *IEEE Sensors Journal*, 19(17):7283–7299, 2019. 2, 3