 https://doi.org/10.24191/jeesr.v29i1.029

SNR-Robust Active Frame Detection for Dual-Radar Fall Monitoring via Maximum Ratio Combining

Nur Iman Syahirah, Nur Emileen Abdul Rashid, Mohd Mustafa Awang Kechik, and Khairul Khaizi Mohd Shariff*

Abstract— **In a practical setting of radar-based fall detection system, radar data arrives as a continuous, stream rather than packaged events. Before the system can recognise a fall, it must first constantly scan this incoming stream to isolate the exact moments of genuine movement from empty background recordings. The conventional approach of using a fixed energy threshold to make this approach vulnerable to failure because it cannot adapt to its unexpected background noise. This situation can easily trigger false alarms, and on the other hand, weak signals from a distant fall might go completely undetected. Furthermore, a threshold perfectly tuned for one room often fails in another, making constant manual recalibration impractical for real-world deployment to address this, we propose a noise-adaptive diversity-combining approach. Based on the principle of maximum-ratio combining, this method weights horizontal and vertical radar energies according to their respective noise levels. We evaluated our approach on a dual-radar dataset against independent time-domain ground truth, testing across simulated signal-to-noise ratios (SNRs) from +20 to −10 dB. As expected, uncalibrated fixed-thresholds performed poorly (mean F1-score: 0.095). However, our proposed method (F1-score: 0.75–0.86) also significantly outperformed calibrated single- and dual-radar constant false-alarm-rate (CFAR) baselines (F1-score: 0.74–0.77) at every SNR level ($p < 0.001$). These results demonstrate that threshold calibration alone is insufficient for robust frame detection, and that our noise-adaptive approach provides consistent improvements regardless of environmental noise.**

Index Terms— **Fall detection, Dual-radar system, CFAR detection, Diversity combining, Active frame detection**

I. INTRODUCTION

According to the World Health Organization (WHO), falls are the second leading cause of death globally. It is estimated 684,000 fatalities cause by fall each year, with over 80% occurring in low- and middle-income countries [1]. Additionally, approximately 37.3 million falls annually are severe enough to require medical attention, affecting roughly 30–40% of older adults worldwide [1]. Delaying help after a fall significantly increases morbidity and mortality, making a timely response absolutely critical. For this reason, the development of automated fall detection systems has become a critical area of technology.

Among the sensing modalities proposed for fall detection, radar (Radio Detection and Ranging) offers several advantages over wearable- and camera-based approaches. Wearable sensors, such as inertial measurement units (IMUs), must be attached to the body and may therefore be inconvenient for continuous use. On the other hand, camera-based systems raise privacy concerns and exhibit degraded performance under poor lighting conditions. Radar mitigates these limitations via contactless sensing without capturing identifiable visual information and operates independently of ambient illumination.

In a practical radar fall detection system, radar data arrives as an unbroken stream rather than as pre-segmented events. Before any higher-level analysis, (e.g., feature extraction and classification) can be performed, the system must first decide which short time-segments of that stream, referred to here as frames, actually correspond to a real physical event (e.g., a person moving) and which contain only background noise, clutter, or an otherwise empty scene. This decision process is known as frame segmentation, and each frame is labeled as either active or judged to contain genuine target-induced motion, or inactive or, judged to contain no useful signal. The reliability of subsequent signal processing depends on this decision being made correctly. If a real event frames are wrongly marked inactive, that portion of the event is effectively discarded before it ever reaches the classifier, and if noise is wrongly marked active, the classifier is instead presented with corrupted or misleading input. It is this frame-level active/inactive decision and how robustly it can be made under varying signal conditions is the focus of this paper.

The existing approach of the active/inactive frame decision is commonly made by comparing the energy of each frame with a predefined threshold [2]. Although computationally simple, this approach is sensitive to variations in signal and noise levels. Changes in the background noise floor may increase false activations, whereas weak target returns caused by distance, or orientation of fall may reduce below the threshold and remain undetected. Consequently, a threshold calibrated for one environment, i.e., room or operating condition may not remain

This manuscript is submitted on July 21, 2026, revised on August 25, 2026, accepted on August 28, 2026 and published on October 31, 2026.

Nur Iman Syahirah, Nur Emileen Abdul Rashid, Khairul Khaizi Mohd Shariff are with the Faculty of Electrical Engineering, Universiti Teknologi MARA, Shah Alam, Selangor, Malaysia.

Mohd Mustafa Awang Kechik is with the Department of Physics, Universiti Putra Malaysia, Serdang, Malaysia.

*Corresponding author
Email address: khairulkhaizi@uitm.edu.my

1985-5389/© 2026 The Authors. Published by UiTM Press. This is an open access article under the CC BY-NC-ND license (http://creativecommons.org/licenses/by-nc-nd/4.0/).

reliable under another. Repeated manual recalibration is impractical for a system deployed across different locations.

These limitations become more pronounced in multi-radar monitoring systems, where simultaneous measurements of the same motion event may exhibit different signal and noise conditions because of sensor orientation and target aspect. Although the complementary observations provide additional information, conventional frame-gating methods typically process each radar independently and combine only the resulting binary decisions. Consequently, potentially useful energy contributions may be discarded before the fusion stage. However, existing studies generally require each radar to detect motion independently before any fusion takes place, rather than incorporating a second sensor's raw signal into the initial active/inactive frame decision. This limitation motivates the present study.

This paper proposes a multi-radar fusion approach for active-frame detection that combines measurements obtained simultaneously from spatially separated radars observing the same motion event. By exploiting the complementary information available across the different sensing viewpoints, the method aims to improve the detection of genuine target-induced motion under varying signal and noise conditions. In particular, it reduces reliance on the response of any single radar and avoids discarding potentially useful information before the fusion stage. We first demonstrate the limitations of conventional fixed-threshold gating when a threshold is transferred unchanged from a different task. Although task-specific recalibration substantially improves its performance, the method remains limited by its equal treatment of the multi radar measurements. The proposed method instead weights each radar's frame energy according to its estimated noise level before fusion, following the principle of maximum-ratio combining used in diversity reception [3].

The objectives of this study are:

- To evaluate the performance limitations of conventional fixed-threshold frame gating under varying SNR conditions
- To propose a noise-weighted diversity-combining method for dual-radar active-frame detection
- To benchmark the proposed method against calibrated CFAR-based fusion baselines using ground-truth labels

The contributions of this work are as follows:

- We demonstrate the limitations of conventional fixed-threshold gating. An uncalibrated threshold yields a mean $FF1$-score of only 0.095, while task-specific calibration improves performance to 0.74–0.77 but remains consistently inferior to the proposed noise-adaptive fusion method.
- A noise-weighted diversity-combining method is proposed, in which the energy from each radar is weighted inversely to its estimated noise level before fusion and thresholding. Across SNRs from +20 to −10 dB, the method significantly outperforms the calibrated baselines in event detection.
- The proposed method is evaluated against fixed-threshold AND and OR fusion, single-radar gating, and local CA-CFAR.

The rest of the paper is organised as follows: Section II provides the related works, Section III explains the methodology, Section IV presents the results and finally Section V provides the conclusion of this work.

II. Related Works

A. Multi-view sensing radar sensing

A fundamental limitation of traditional single-radar sensing is that it can only measure radial velocity, which is motion traveling directly toward or away from the sensor. Because a single radar cannot directly observe tangential (sideways) motion, its performance drops drastically when a target moves perpendicular to its line of sight [4]. To overcome this orientation 'blind spot' researchers have increasingly turned to spatially diverse, multi-view radar configurations.

In these multi-view setups, the standard practice is to place radars at different angles such as a dual-radar layout [5] where two sensors are positioned perpendicular to one another, to ensure that if one radar is blind to a movement, the second radar captures it. Current systems typically fuse the data from these different viewpoints at the later stages of the processing pipeline. For instance, some approaches extract motion features from each radar independently and combine them afterward to improve classification [6], while others let each radar make its own independent guess before using a voting-based or neural network fusion method to make a final decision [7], [8].

However, all of these existing multi-view approaches share a unaddressed issue. They assume that each radar is capable of detecting whether motion is occurring entirely on its own, prior to any data sharing. In other words, the initial decision of which time-segments contain usable target motion is made independently by each sensor using only its own signal statistics. No existing method utilises the second radar's viewpoint during this initial detection stage. Consequently, if a single radar suffers from high noise or orientation blindness and fails to register a frame on its own, that crucial information is permanently discarded before the multi-view fusion step ever has a chance to occur.

B. Temporal segmentation and learned frame classifiers

An alternative to frame-wise motion detection is to formulate the task as a temporal segmentation or change-point detection problem, analysing the temporal evolution of the signal rather than thresholding each frame independently. Classical techniques such as cumulative-sum statistics and likelihood-ratio tests, as well as learning-based models such as bidirectional long short-term memory networks (Bi-LSTM) applied to radar spectrograms, have been used to partition continuous sensor streams into activity segments [9],[10].[11].

These methods can achieve high segmentation accuracy, but their performance depends strongly on the representativeness of labelled training data, and, more importantly for this work, they

are formulated using the signal stream from a single radar

Consequently, when the response of the first radar is degraded by low SNR, an unfavourable target aspect, or weak target-induced modulation, the onset and termination of an activity may therefore remain difficult to identify. Critically, none of these segmentation approaches, classical or learned, incorporate a second radar's viewpoint into the initial frame-level decision itself.

C. Active frame detection and adaptive thresholding

The problem of deciding whether a given frame contains a genuine target return, as opposed to noise or clutter, is most commonly addressed through energy thresholding. The simplest form uses a fixed threshold, set as a constant multiple of an estimated background noise level [2]. A more principled alternative, and the standard approach in classical radar target detection, is constant false-alarm-rate (CFAR) processing [12]. CFAR estimates the noise level locally, from a window of neighboring frames, and sets the detection threshold relative to that local estimate rather than using one fixed threshold for an entire recording. This allows the threshold to adapt as the background noise level varies across a recording. Many refinements of CFAR for handling non-uniform clutter and multiple-target conditions has been investigated [13], [14]. However, CFAR and its variants share the same underlying property as fixed thresholding. The detection decision for a given radar is derived entirely from that radar's own signal. To date, no existing work has incorporated a second, independently-viewing sensor as part of the detection statistic itself.

III. METHODOLOGY

This section introduces the signal preprocessing for active frame detection. Subsequently, the proposed maximum ratio combining method is presented followed by CFAR baselines. Next, the hyperparameters and ground truth are explained. Finally, the dataset used in this work is explained in detail. The complete workflow of the proposed system is illustrated in the flowchart in Fig. 1.

A. Dataset

This study evaluates the proposed gating method using an existing dual-radar dataset previously acquired at the university. The dataset contains 3,620 events of human motions recorded from 10 participants. The recorded actions includes falls and non-fall activities of daily living, such as sitting, standing, walking, and bending. Each motion event are recorded using an orthogonal radar arrangement.As shown in Fig. 2, one radar was positioned horizontally at a height of 1.3 meters, while the second was mounted 3 meters above the ground and directed downward, with a 3 m baseline distance between their mounting positions. Each recorded event contains four discrete in-phase and quadrature channels: I_H and Q_H from the horizontally oriented radar, and I_V and Q_V from the vertically oriented radar.

The data were acquired using two RF Beam GmbH K-LC2 continuous-wave (CW) radars operating at a center frequency of 24 GHz. The radar modules transmit with an EIRP of +15 dBm with noise floor of approximately −109 dBc. At the short distances used in this environment, human body reflections at the receiver yield with a median of approximately 20 dB across a representative random sample of events.

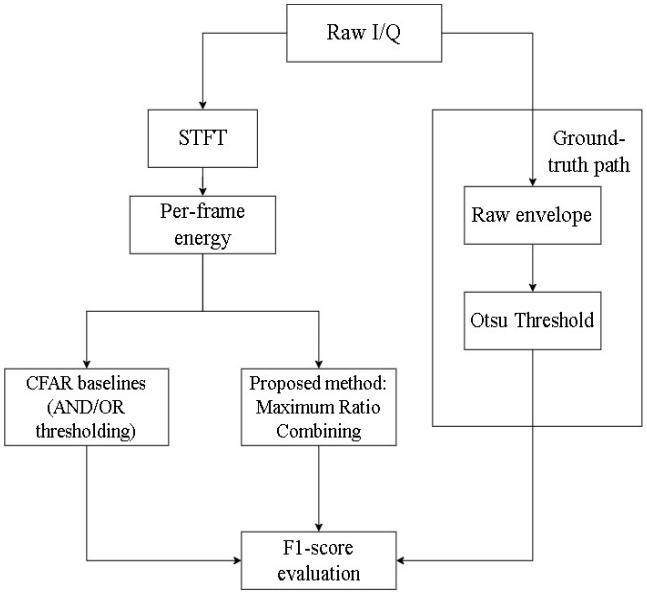

Fig. 1. Flowchart of project.

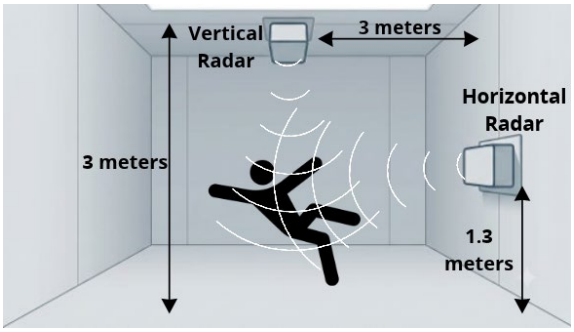

Fig. 2. Illustration of dual-radar system for human activity detection.

B. Preprocessing

Each event was recorded as four baseband channels: in-phase and quadrature signals from the horizontal radar (I_H, Q_H) and from the vertical radar (I_V, Q_V). For each radar $i \in H, V$, the complex baseband signal is formed as

$$s_i[t] = I_i[t] + j\, Q_i[t] \tag{1}$$

A time-frequency representation is obtained for each radar via the short-time Fourier transform (STFT), applied with a window length of L = 1024 samples and overlap of O = 896 samples to produces a spectrogram $Z_i[f, n]$ indexed by frequency bin f and frame index $n = 1, \ldots, N$, where N is the total number of frames per event. To suppress zero-Doppler clutter arising from static reflections, the spectrogram magnitude is masked within a small band of b = 3 bins centered on zero Doppler.

$$|\hat{Z}_i[\mathrm{f,n}]| = \begin{cases} 0, & |f| \le b \\ |Z_i[f,n]|, & \text{otherwise} \end{cases} \quad (2)$$

The per-frame energy for radar i is then defined as the sum of squared masked magnitude across all frequency bins.

$$e_i[n] = \sum_f |\hat{Z}_i[\mathrm{f,n}]|^2 \quad (3)$$

The sequences $e_H[n]$ and $e_V[n]$ computed from Equation (3) are the basis for all gating methods described in the remainder of this section.

C. CFAR methods

Four baseline detection rules are evaluated against the proposed method, all belonging to the constant-false-alarm-rate (CFAR) family. Each derives its detection threshold from an estimate of the local or event-wide noise level, but unlike the proposed method makes its decision using only a single radar statistics, or combines two radars independent single-radar decisions only at the final decision level. For a given energy sequence e and sensitivity multiplier k, a per-event adaptive threshold is defined as

$$\tau(e,k) = \mathrm{median}(e) + k \cdot \mathrm{std}(e) \quad (4)$$

1) Baseline 1: Fixed CFAR-AND

In this approach, a frame is marked active only if both radars independently exceed this threshold given by [15]

$$g_{AND}[n] = \mathbb{1}[e_H[n] > \tau(e_H, k_{AND})] \wedge \mathbb{1}[e_V[n] > \tau(e_V, k_{AND})] \quad (5)$$

In this work, we evaluate two variants of this rule. The first is a naive variant $k_{AND} = 3.0$ which is inherited unchanged from the single-radar classification pipeline for which it was originally developed. The second is a calibrated variant, where the threshold k_{AND} is tuned specifically for the current task (detailed in Section D). Comparing these two approaches isolates the specific impact of calibration, independent of any changes to the underlying detection logic.

2) Baseline 2: Fixed CFAR-OR

This approach has an identical threshold rule, but combined with a logical OR rather than AND, so that either radar alone is sufficient. It is given as [15]

$$g_{OR}[n] = \mathbb{1}[e_H[n] > \tau(e_H, k_{OR})] \vee \mathbb{1}[e_V[n] > \tau(e_V, k_{OR})] \quad (6)$$

Between these first two baseline, CFAR-AND declares a frame active only when both radars independently detect motion, whereas CFAR-OR declares it active when either radar detects motion. CFAR-AND is more conservative, while CFAR-OR is more sensitive to motion that appears strongly in only one radar.

3) Baseline 3: Single-radar CFAR

This approach only use the horizontal radar energy measurements. Contribution from the vertical radar is ignored.

$$g_{single}[n] = \mathbb{1}\big[e_H[n] > \tau\big(e_H, k_{single}\big)\big] \quad (7)$$

4) Baseline 4: Local frame-averaging CFAR (CA-CFAR)

Rather than applying a single threshold derived from the entire event, CA-CFAR estimates the local noise level for each frame under test (FUT) using a sliding window of $2R + 1$frames. A guard region of $2G + 1$ frames centred on the FUT is excluded to prevent target energy from biasing the noise estimate. The resulting detection threshold therefore adapts to local variations in the noise floor is given by [16]

$$\tau_{\mathrm{loc}}[n] = \mu(e[R(n) \setminus G(n)]) + k_{\mathrm{ca}} \cdot \mathrm{std}(e[R(n) \setminus G(n)]) \quad (8)$$

where, μ is the mean, $R(n) = [n - R, n + R]$ denotes the reference window centred on the nth frame under test, $G(n) = [n - G, n + G]$ denotes the excluded guard region, and k_{ca} is the CA-CFAR sensitivity multiplier. CA-CFAR is applied independently to the two radar signals, and the resulting binary decisions are combined using the AND rule as

$$g_{\mathrm{ca}}[n] = \mathbb{1}\big[e_H[n] > \tau_{\mathrm{loc}}(e_H)[n]\big] \wedge \mathbb{1}\big[e_V[n] > \tau_{\mathrm{loc}}(e_V)[n]\big] \quad (9)$$

D. Maximum ratio combining

The calibrated CFAR baselines fuse the radar outputs only after each radar has independently classified a frame as active or inactive using its own energy measure. The resulting binary decisions are then combined using an AND or OR rule. This approach discards information on the degree to which each radar's energy exceeds its threshold and treats both radars as equally reliable, regardless of their noise levels.

In contrast, the proposed method forms a single detection statistic by weighting each radar frame energy according to its estimated reliability. This principle is analogous to maximum-ratio combining in diversity reception, where noisier branches are assigned lower weights in the fused statistic [3]. For each radar i in {H,V}, a per-event noise-level estimate is obtained from the median of that radar own frame energy:

$$n_i = \mathrm{median}(e_i) \quad (10)$$

Each radar combining weight is then set inversely proportional to its own noise-level estimate, so that a radar with a higher noise floor contributes less to the combined statistic

$$w_i = \frac{1}{n_i + \varepsilon} \quad (11)$$

where ε is a small constant (10^{-30} in this work) preventing division by zero. The combined detection statistic for frame n is the noise-weighted sum of both radars frame energy

$$e_{comb}[n] = w_H\, e_H[n] + w_V\, e_V[n] \quad (12)$$

A single adaptive threshold, of the same form as Equation (4), is then applied to this combined statistic, using a sensitivity multiplier k_{ws} specific to the weighted-sum method

$$g_{ws}[n] = \mathbb{1}[e_{comb}[n] > \tau(e_{comb}, k_{ws})] \quad (13)$$

Only frames for which the decision output equals 1 are forwarded to the feature-extraction stage, while all others are treated as inactive according to the established convention. Two key properties distinguish this combining rule from the CFAR baselines. First, because the energy from both radars is fused into a single statistic prior to thresholding, a frame can still be correctly marked as active even if neither radar individually exceeds its own threshold. This occurs when their joint sub-threshold contributions surpass the combined threshold valuable information that is entirely discarded by decision-level combinations (e.g., AND or OR logic). Second, weighting each radar's contribution by its respective noise level ensures that a sensor with an unfavourable orientation or a locally elevated noise floor is automatically down-weighted. This adaptive approach does not require prior knowledge of the noise level or manual recalibration, as the weight in Equation (11) is independently re-estimated for every event using its own data.

A curve-fitting approach to the noise floor would require assuming a specific noise model (e.g., a fixed decay or parametric clutter distribution) that may not hold across rooms, radar orientations, or target aspects. Explicit clutter modelling, similarly, requires scene-specific calibration that is impractical for a system deployed across unseen environments. In contrast, the proposed weighting Equation (11) requires only a per-event median energy estimate from each radar, making it self-calibrating without any assumption about the underlying noise distribution or clutter structure.

E. Hyperparameters tuning

All methods with tunable sensitivity parameters were calibrated using a held-out tuning subset by grid search maximising mean F_1-score against the independent ground truth across representative SNR conditions (+20 dB and −10 dB).

The grid search encompassed three distinct algorithmic approaches, each with its own defined parameter space

- *Fixed CFAR Baselines*: The calibrated sensitivity multipliers k_{AND}, k_{OR}, and k_{single} were selected independently from the grid {0.1, 0.25, 0.5, 0.75, 1.0, 1.5, 2.0, 3.0}.
- CA-CFAR baseline: The search grid for the reference-window half-width R, was selected from {5, 10, 15, 20}, G in {1, 2, 3}, and k_{ca} in {0.25, 0.5, 0.75, 1.0, 1.5, 2.0}.
- Maximum-ratio combining method, for which the sensitivity multiplier k_{ws} was selected from the same grid used for the fixed baselines.

The final selected values are as follows:

- *Fixed CFAR Baselines*: k_{AND} = 0.25, k_{OR} = 1.0, k_{single} = 0.25
- CA-CFAR baseline: R = 20, G = 5, k_{ca} = 0.5
- Maximum-ratio combining: k_{ws} = 0.75.

Notably, the fixed CFAR-AND baseline was deliberately excluded from this calibration procedure. It retains its original value of k_{AND_naive} = 3.0 unchanged from the single-radar processing pipeline it was originally developed. This omission was intentional as it serves to demonstrate the performance consequences of transferring a pre-set threshold across different tasks without proper recalibration.

F. Independent ground truth

Evaluating these methods requires frame-level ground-truth labels. Deriving these labels from the same STFT frame energy used by the evaluated detectors risks circular evaluation, as it would inherently bias the results toward methods with similar decision rules. To ensure a fair evaluation, our ground truth is derived from the raw time-domain amplitude envelope instead of time-frequency domain.

For each radar i in {H,V}, the raw complex baseband signal $s_i[t]$ forms a time-domain amplitude envelope

$$env_i[t] = |\, s_i[t] - \mathrm{mean}(s_i)\,| \quad (14)$$

The two radars envelopes are summed, and averaged within the time window corresponding to each STFT frame n to produce a per-frame time-domain activity statistic $v[n]$

$$v[n] = \mathrm{mean}(env_H[t] + env_V[t]) \quad (15)$$

Otsu's method [17] is first applied to $v[n]$ to determine the threshold τ_{otsu} that maximizes the between-class variance of this time-domain statistic. Frames exceeding this threshold are assigned to the active class. To remove isolated detections and fill short gaps, this sequence is refined using the exact same operations applied in the STFT-domain by closing with a three-frame structuring element followed by opening with a two-frame structuring element. This refinement yields the final ground-truth sequence

$$g_T[n] = \mathrm{open}(\mathrm{close}(\mathbb{1}[v[n] > \tau_{Otsu}])) \quad (16)$$

To ensure a consistent evaluation of every gating method, this ground-truth sequence is generated only once per event using clean, non-degraded signals. The resulting labels are then held constant across all simulated SNR conditions.

$$\mathrm{SNR}_{\mathrm{dB}} = 10\log_{10}\left(\frac{P_{s_i}}{P_{n_i}}\right), \qquad P_{s_i} = \frac{1}{T}\sum_t |s_i[t]|^2 \quad (17)$$

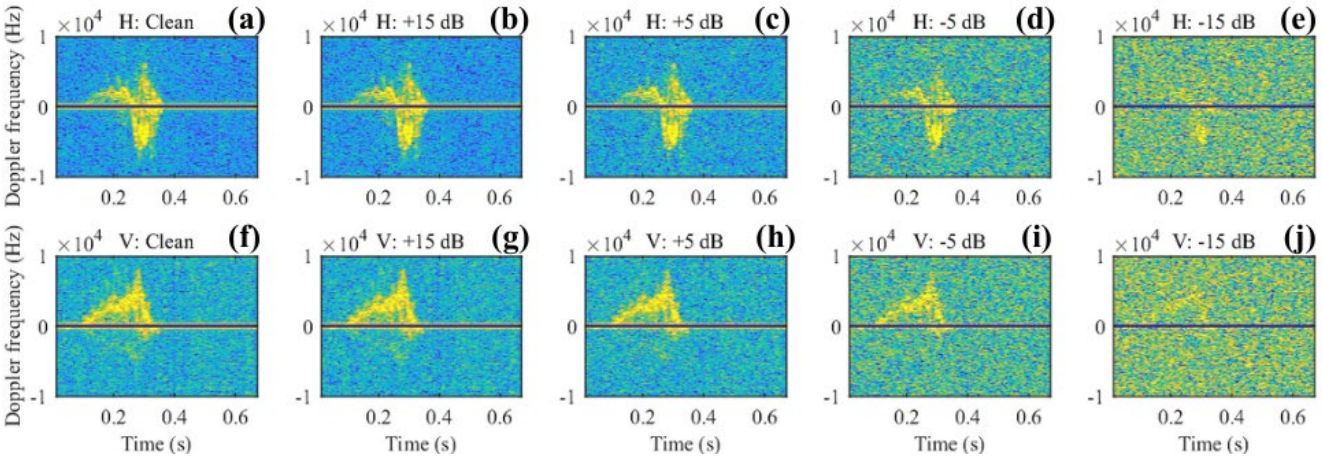

Fig. 3. Spectrograms of a fall event from the horizontal (a–e) and vertical (f–j) radars across decreasing SNR: clean, +15 dB, +5 dB, −5 dB, −15 Db.

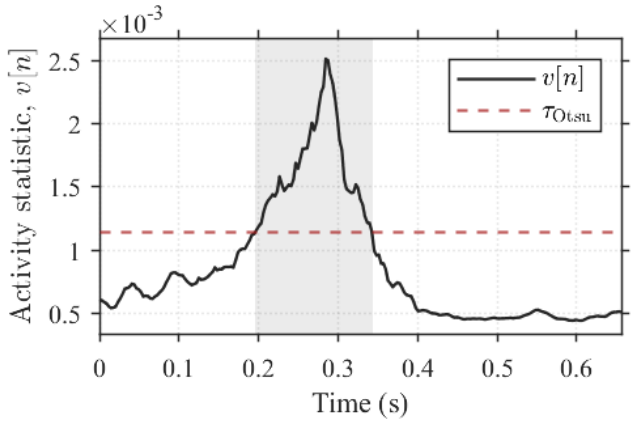

Fig. 4. Independent time-domain ground-truth construction for the event in Fig 3. Because the activity statistic $v[n]$ is derived from the raw amplitude envelope rather than the STFT energy used by the evaluated detectors, this confirms the ground truth is not circular with respect to the methods being tested.

Simulated SNR degradation was used rather than re-recording at multiple physical noise levels because it allows the exact same underlying event and therefore the exact same ground-truth labels to be evaluated identically across every SNR condition. Re-acquiring physical data at controlled SNR levels is impractical to control precisely (ambient RF noise, receiver gain drift) and would confound noise-level effects with unavoidable differences between separately recorded events

The degraded signal for the ith radar is modelled as $s_i'[t] = s_i[t] + n_i[t]$, where $n_i[t]$ denotes an independently generated noise realization for each radar at each SNR level. Prior to the AWGN added, DC offset signal is removed first.

For each event, gating method, and SNR level, the predicted active-frame mask is compared against ground truth using standard detection metrics namely precision, recall and F_1-score.

G. *Implementation*

All signal processing and evaluation were implemented in MATLAB R2021a. The CFAR baselines and the proposed MRC weighting were both implemented and evaluated within this same environment, keeping the entire processing pipeline consistent and reproducible across all reported experiments. The CA-CFAR baseline was implemented following the standard cell-averaging CFAR algorithm as described in [16], using the parameters specified in Section III-E.

IV. RESULTS

A. *Signal characteristics and baseline validation*

Before quantifying the performance of the proposed and baseline detection methods, it is necessary to establish the visual impact of the simulated noise on the radar signals and validate the reliability of our independent ground truth. Fig 3 The horizontal radar spectrograms are shown in panels (a)–(e) and the vertical radar spectrograms in panels (f)–(j), under the clean condition and four levels of simulated SNR degradation, respectively. By observing how the characteristic Doppler signature blurs at +5 dB (panels (c), (h)) and becomes entirely indistinguishable from the noise floor by −15 dB (panels (e), (j)), we can visually establish the challenge that fixed-threshold detectors face under varying signal conditions.

Furthermore, to ensure a fair evaluation, we demonstrate that our time-domain ground truth accurately captures these events without relying on the time-frequency methods being evaluated. Fig 3 plots the time-domain activity statistic for the same event, demonstrating that our procedure successfully isolates a single, temporally contiguous active interval that aligns with the visual event structure in Fig 3, (e.g., panels (a) and (f)), rather than fragmenting into multiple disconnected detections.

Otsu's method (Equation 16) selects the threshold $\tau_{Otsu} \approx 1.15 \times 10^{-3}$(dashed red line) that separates this elevated region from the surrounding baseline. The shaded region marks the resulting active-frame interval after the closing and opening are applied to remove isolated detections and bridge short gaps.

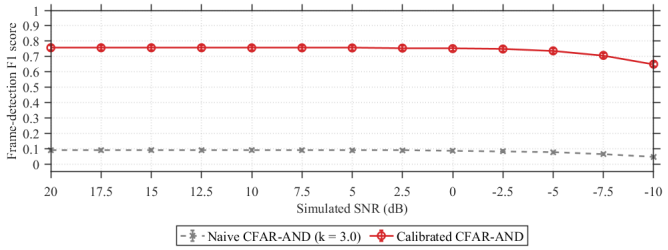

Fig. 5. Performance comparison of naive vs. calibrated CFAR-AND baselines.

TABLE I. PERFORMANCE COMPARISON OF NAÏVE VS. CALIBRATED CFAR-AND

Method	+20db	+10db	0db	-10db
Naïve CFAR-AND	0.09	0.09	0.09	0.05
Calibrated CFAR-AND	0.75	0.75	0.74	0.65

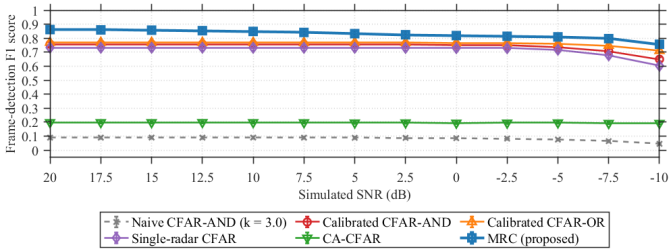

Fig. 6. Performance comparison All CFARs and the proposed method.

TABLE II. PERFORMANCE COMPARISON ACROSS ALL BASELINES AND THE PROPOSED METHOD

Method	+20dB	+10dB	0dB	-10dB
Naïve CFAR-AND	0.09	0.09	0.09	0.04
Calibrated CFAR-AND	0.75	0.75	0.74	0.65
Calibrated CFAR-OR	0.77	0.77	0.77	0.70
Single-radar CFAR	0.73	0.73	0.73	0.60
CA-CFAR	0.20	0.20	0.19	0.18
MRC (Proposed)	0.86	0.85	0.85	0.75

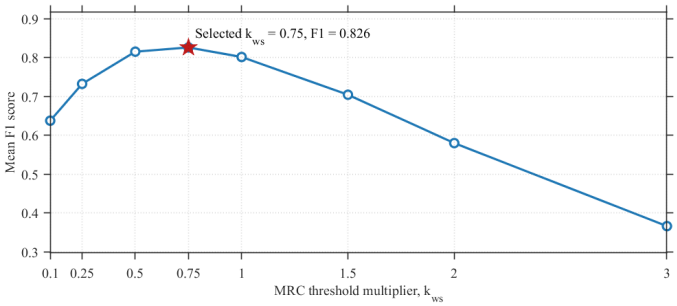

Fig. 7. Hyperparameter sensitivity of the proposed MRC method.

B. Performance comparison of naive CFAR-AND versus calibrated CFAR-AND

To establish the necessity of task-specific calibration before evaluating different fusion strategies, we first show the impact of threshold calibration using the Fixed CFAR-AND baseline. Fig. 5 demonstrates that transferring a naive threshold from a previous single-radar task results in uniform failure, yielding a mean F_1-score of only approximately 0.09 at high SNR. Because this poor performance remains flat across the entire +20 dB to −10 dB sweep, it proves that the naive threshold fails regardless of signal quality simply because its sensitivity was never matched to this task

This result provides the main motivation for Contribution 1. In dual-radar active-frame detection, task-specific threshold calibration has a greater influence on performance than the subsequent fusion rule, whether AND, OR, or diversity combining. Meaningful comparisons among detection methods therefore require all thresholds to be calibrated under the same evaluation conditions. The large performance gap observed here supports using the naive variant only as a reference rather than as a competitive baseline.

C. Performance comparison of proposed method vs all baselines

Fig. 6 compares the F_1 scores of all six methods across simulated SNR levels from +20 to −10 dB. The results form three clear performance groups. The proposed maximum-ratio combining (MRC) method consistently achieves the highest F_1 score at every SNR level, decreasing from approximately 0.86 at +20 dB to 0.75 at −10 dB. Its improvement over the strongest baseline decreases from about 0.09 at high SNR to 0.04 at −10 dB, but remains evident throughout the evaluated range.

This result indicates that fusing the radar energies into a single reliability-weighted statistic before thresholding is more effective than combining independent binary decisions from the two radars.

The calibrated CFAR-OR, CFAR-AND, and single-radar CFAR methods form the intermediate performance group. From +20 to approximately −2.5 dB, these methods achieve similar F_1 scores of about 0.75–0.77. Their performance begins to diverge below 0 dB. This difference suggests that the AND rule becomes increasingly restrictive when one radar is strongly affected by noise, while the OR rule remains more tolerant of unequal signal quality across the two sensors. The single-radar CFAR method degrades more rapidly and reaches approximately 0.60 at −10 dB which demonstrates the limitation of relying on a single sensing viewpoint under low-SNR conditions. This degradation is consistent with the orientation-blindness limitation described for single-viewpoint radar sensing [4] and confirms empirically that the second radar's viewpoint carries information the single-radar baseline cannot recover even under favorable SNR.

The remaining methods exhibit substantially lower performance. The calibrated CA-CFAR and the naive CFAR-AND baseline perform worst. This confirms that the use of an uncalibrated threshold causes systematic detection failure rather than degradation driven primarily by SNR.

D. Hyperparameter sensitivity

Fig. 7 shows the sensitivity of the proposed maximum-ratio combining (MRC) method to the threshold multiplier k_{ws}. The parameter was evaluated on a held-out tuning subset using the mean F_1 score across SNR +20 and (−10) dB. The results exhibit a clear single maximum.

This trend reflects the tradeoff introduced by the adaptive threshold. Small values of k_{ws} produce a permissive decision

boundary, increasing the likelihood that noise-dominated frames are classified as active. In contrast, excessively large values produce a restrictive threshold that suppresses low-amplitude target responses and increases missed detections. The pronounced dependence on k_{ws} demonstrates the importance of task-specific calibration. Accordingly, (k_{ws} = 0.75, which achieved the highest mean F_1 score and was selected for all subsequent evaluations.

V. CONCLUSION

This study proposed a noise-adaptive fusion method for active-frame detection in dual-radar fall monitoring by combining radar energies before thresholding. The proposed method consistently achieved the highest F1-score across SNR levels from +20 to −10 dB, ranging from 0.75 to 0.86, compared with 0.70 to 0.77 for the strongest calibrated baseline. The results show that reliability-weighted pre-threshold fusion can improve active-frame detection when the two radar observations experience different signal and noise conditions.

The method requires only an estimate of the radar noise level and does not rely on a predefined noise or clutter model. Future work will validate the approach under different radar geometries and physical environments, implement real-time noise estimation, and evaluate its impact on end-to-end fall classification.

ACKNOWLEDGMENT

The authors gratefully acknowledge members of Applied Electromagnetics Research Group (AERG) for the development of the dual-radar system and its dataset, and the Universiti Teknologi MARA (UiTM) for providing the facilities and institutional support required to conduct this work

REFERENCES

[1] World Health Organisation, "Falls," World Health Organisation. Accessed: Jul. 20, 2026. [Online]. Available: https://www.who.int/news-room/fact-sheets/detail/falls

[2] F. Savvidou, S. A. Tegos, P. D. Diamantoulakis, and G. K. Karagiannidis, "Passive Radar Sensing for Human Activity Recognition: A Survey," *IEEE Open Journal of Engineering in Medicine and Biology*, vol. 5, pp. 700–706, 2024.

[3] D. G. Brennan, "Linear diversity combining techniques," *Proceedings of the IEEE*, vol. 91, no. 2, pp. 331–356, Feb. 2003.

[4] Y. Yang, C. Hou, Y. Lang, T. Sakamoto, Y. He, and W. Xiang, "Omnidirectional Motion Classification With Monostatic Radar System Using Micro-Doppler Signatures," *IEEE Transactions on Geoscience and Remote Sensing*, vol. 58, no. 5, pp. 3574–3587, May 2020.

[5] K. Saho, S. Hayashi, M. Tsuyama, L. Meng, and M. Masugi, "Machine Learning-Based Classification of Human Behaviors and Falls in Restroom via Dual Doppler Radar Measurements," *Sensors*, vol. 22, no. 5, p. 1721, Jan. 2022.

[6] X. Feng *et al.*, "Millimeter-Wave Radar Monitoring for Elder's Fall Based on Multi-View Parameter Fusion Estimation and Recognition," *Remote Sensing*, vol. 15, no. 8, p. 2101, Jan. 2023.

[7] M. Ritchie, F. Fioranelli, H. Borrion, and H. Griffiths, "Multistatic micro-Doppler radar feature extraction for classification of unloaded/loaded micro-drones," *IET Radar, Sonar & Navigation*, vol. 11, no. 1, pp. 116–124, 2017.

[8] R. G. Guendel, F. Fioranelli, and A. Yarovoy, "Distributed radar fusion and recurrent networks for classification of continuous human activities," *IET Radar Sonar Navig.*, vol. 16, no. 7, pp. 1144–1161, 2022.

[9] S. Aminikhanghahi and D. J. Cook, "A survey of methods for time series change point detection," *Knowl Inf Syst*, vol. 51, no. 2, pp. 339–367, May 2017.

[10] S. Aminikhanghahi, T. Wang, and D. J. Cook, "Real-Time Change Point Detection with Application to Smart Home Time Series Data," *IEEE Transactions on Knowledge and Data Engineering*, vol. 31, no. 5, pp. 1010–1023, May 2019.

[11] A. Shrestha, H. Li, J. Le Kernec, and F. Fioranelli, "Continuous Human Activity Classification From FMCW Radar With Bi-LSTM Networks," *IEEE Sensors Journal*, vol. 20, no. 22, pp. 13607–13619, Nov. 2020.

[12] V. G. Hansen and H. R. Ward, "Detection Performance of the Cell Averaging LOG/CFAR Receiver," *IEEE Transactions on Aerospace and Electronic Systems*, vol. AES-8, no. 5, pp. 648–652, Sep. 1972, doi: 10.1109/TAES.1972.309580.

[13] M. Y. Rihan, Z. B. Nossair, and R. I. Mubarak, "An improved CFAR algorithm for multiple environmental conditions," *SIViP*, vol. 18, no. 4, pp. 3383–3393, Jun. 2024.

[14] W. Zhang, Y. Li, Z. Zheng, L. Xu, and Z. Wang, "Multi-Target CFAR Detection Method for HF Over-The-Horizon Radar Based on Target Sparse Constraint in Weibull Clutter Background," *Remote Sens.*, vol. 15, no. 10, 2023.

[15] P. K. Varshney, *Distributed Detection and Data Fusion*. New York, NY: Springer, 1997.

[16] C. R. Barrett, "Adaptive Thresholding and Automatic Detection," in *Principles of Modern Radar*, J. L. Eaves and E. K. Reedy, Eds., Boston, MA: Springer US, 1987, pp. 368–393.

[17] N. Otsu, "A Threshold Selection Method from Gray-Level Histograms," *IEEE Transactions on Systems, Man, and Cybernetics*, vol. 9, no. 1, pp. 62–66, Jan. 1979.

.