A variance-reduced simulation for benchmarking CFAR detectors in non-homogeneous environments

Pham Van Hung[1*], Tran Thi Nhung[2]

[1]Le Quy Don Technical University, 236 Hoang Quoc Viet, Nghia Do, Hanoi, Vietnam;
[2]Department of Electrical and Electronic Engineering, Nam Dinh University of Technology and Education, 186 Phu Nghia, Thien Truong, Ninh Binh.
[*]Corresponding author: hungpv_k31@lqdtu.edu.vn
Received 12 Jan. 2026; Revised 29 Apr. 2026; Accepted 11 May 2026; Published 25May 2026.
DOI: https://doi.org/10.54939/1859-1043.j.mst.111.2026.39-51

ABSTRACT

This paper presents an accurate and stable simulation framework for evaluating the detection performance of constant false alarm rate (CFAR) detectors in the presence of clutter non-homogeneity and interference. The proposed framework incorporates theoretically consistent threshold calibration for CFAR detectors and introduces a variance-reduced estimator for false-alarm probability based on conditional expectation. This approach enables reliable estimation of extremely small false-alarm probabilities without excessive computational cost. Clutter edge scenarios are modeled using an edge-aware formulation that ensures physical consistency between the reference window and the cell under test. Extensive numerical results validate the proposed framework under homogeneous clutter and demonstrate its effectiveness in analyzing detector robustness in the presence of multiple interfering targets and clutter edges. The results provide clear insight into the relative strengths and limitations of different CFAR detectors and highlight the importance of accurate simulation techniques for meaningful performance comparison in realistic radar environments.

Keywords: CFAR; Trimmed-mean; Variability-index; Monte Carlo; Variance reduction.

1. INTRODUCTION

Constant-false-alarm-rate (CFAR) processing remains a core component of radar target detection because it maintains a prescribed false-alarm probability under uncertain and time-varying clutter power [1, 2]. In the classical sliding-window setting with independent exponentially distributed clutter samples, a broad family of CFAR detectors admits closed-form (or numerically invertible) threshold calibrations for a desired probability of false alarm (P_{fa}) [3, 4]. However, practical operation is dominated by non-homogeneous effects such as multiple interfering targets in the reference window and clutter-power transitions (clutter edges), which can cause severe false-alarm mis-regulation and target masking [3, 5, 6].

A number of robust CFAR variants have therefore been widely adopted. The cell-averaging (CA) CFAR is optimal in homogeneous exponential clutter but is sensitive to interference [4]. The greatest-of (GO) and smallest-of (SO) detectors improve robustness near clutter edges at the expense of detection loss or vulnerability in multi-target situations [3, 6]. Ordered-statistic (OS) and trimmed-mean (TM) detectors mitigate multi-target contamination through censoring/ ordering of the reference samples [7−12]. Finally, variability-index (VI) CFAR performs a data-driven switching to select an appropriate reference statistic based on local heterogeneity [13].

Despite mature analytical theory, simulation is still indispensable for (i) validating threshold calibration, (ii) quantifying detection/false-alarm behavior under composite operating conditions, and (iii) supporting reproducible comparisons across detectors [14−18]. A key difficulty is that direct Monte Carlo estimation of rare-event P_{fa} (e.g., $10^{-6} \div 10^{-10}$) requires prohibitively many trials and yields unstable curves with large uncertainty bands [19]. Recent advances in rare-event simulation and variance reduction techniques have sought to address these computational

bottlenecks by leveraging more efficient sampling strategies [20, 21].

While extensive theoretical derivations for individual CFAR detectors exist, a rigorous and standardized benchmarking tool capable of evaluating these algorithms at extremely low false-alarm regimes ($P_{fa} < 10^{-6}$) under complex scenarios remains lacking due to computational constraints. Motivated by this gap, rather than proposing a new detection algorithm, this paper contributes a unified, variance-reduced simulation framework for benchmarking existing CFAR variants. Specifically, the contributions are threefold:

(i) we implement theory-consistent threshold calibration for each detector using established closed-form or numerically invertible P_{fa} expressions [3, 4, 7];

(ii) we employ a variance-reduced estimator based on conditional expectation to enable accurate P_{fa} evaluation in the rare-event regime together with rigorous confidence bands [19, 22];

(iii) we provide systematic performance studies for multi-target interference and clutter-edge transitions, illustrating detector robustness and target masking mechanisms [3, 12, 18].

The remainder of this paper is organized as follows. Section 2 introduces the signal and detection model, and reviews the CFAR detectors considered in this study and their threshold calibration. Section 3 describes the proposed variance-reduced simulation methodology. The simulation setup and numerical results are presented and discussed in Section 4, and concluding remarks are given in Section 5.

2. PROBLEM

2.1. Signal and clutter models

Consider a pulsed radar system employing square-law detection and a CFAR model, where detection is performed independently in each range cell, as shown in Figure 1. Let Y denote the output of the cell under test (CUT), and let X_i (i=1,...,N) denote the outputs of the surrounding reference cells used for adaptive threshold estimation.

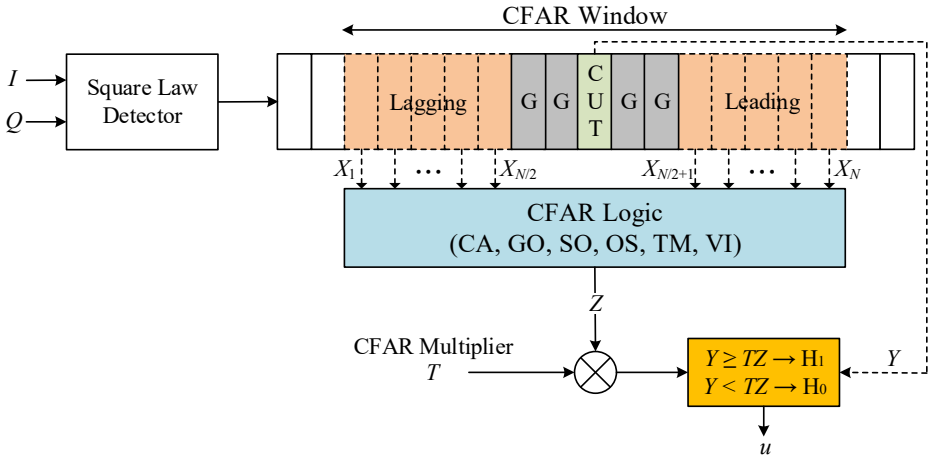

Figure 1. *Radar CFAR detection.*

Under the null hypothesis H_0 (absence of a target in the CUT), the received signal is assumed to consist of clutter and thermal noise only. In this work, clutter is modeled as a complex Gaussian random process, such that the square-law detector output follows an exponential distribution. Specifically,

$$Y \sim \mathrm{Exp}(\mu_{\mathrm{CUT}}), \qquad X_i \sim \mathrm{Exp}(\mu_i) \tag{1}$$

where μ_{CUT} and μ_i denote the mean clutter powers in the CUT and the i-th reference cell, respectively. Under homogeneous conditions, all reference cells and the CUT share the same mean,

i.e., $\mu_{CUT} = \mu_i = 1$. In non-homogeneous environments, such as clutter edges or multi-target scenarios, the means μ_i may vary across the reference window, while the CUT mean is determined by its position relative to the clutter transition.

Under the alternative hypothesis H_1, the CUT contains a target with a signal-to-noise ratio (SNR), e.g. ρ, and the detector output is modeled as

$$Y \sim \text{Exp}(1+\rho) \tag{2}$$

assuming a nonfluctuating target model. Interfering targets within the reference window are modeled similarly by increasing the corresponding mean values μ_i.

This statistical model is widely adopted in the CFAR literature and provides a tractable yet physically meaningful basis for performance analysis and simulation. While more complex statistical models, such as non-Gaussian clutter (Weibull or Compound-Gaussian distributions) [23, 24] and Swerling fluctuating targets, are highly relevant in modern radar scenarios, the exponential model provides a tractable and widely accepted baseline. Using this baseline allows us to rigorously validate the variance-reduction methodology and threshold calibration mechanisms against known analytical closed forms, which is the primary focus of this paper.

2.2. CFAR detectors and threshold calibration

The general CFAR detection rule compares the CUT output Y to an adaptive threshold $\tau(X)$ derived from the reference cells:

$$Y \underset{H_0}{\overset{H_1}{\gtrless}} \tau(X) \tag{3}$$

The threshold $\tau(X)$ is constructed such that the P_{fa} remains constant and equal to a predefined design value $P_{fa,0}$, independent of the unknown clutter power under homogeneous conditions. A key property of CFAR detectors is scale invariance, whereby scaling all reference samples and the CUT by the same factor does not alter the P_{fa}.

Different CFAR schemes differ in the functional form of $\tau(X)$, depending on how the reference data are processed. The following content briefly reviews the CFAR detectors considered in this study and discusses their threshold calibration.

2.2.1. CA-CFAR

Let $S = \sum_{i=1}^{N} X_i$. The CA rule is $Y > C_N S$. For i.i.d. exponential clutter,

$$P_{fa}^{\text{CA}} = (1+C_N)^{-N}, \qquad P_d^{\text{CA}}(\rho) = (1+C_N/(1+\rho))^{-N} \tag{4}$$

hence

$$C_N = P_{fa,0}^{-1/N} - 1 \tag{5}$$

In a nonhomogeneous reference window where $X_i \sim \text{Exp}(\text{mean} = \mu_i)$ are independent but not identically distributed (e.g., $M_{I/C}$ interference/clutter-edge cells with $\mu_i = \beta$ and the rest $\mu_i = 1$), the Laplace product yields

$$P_{fa}^{\text{CA}} = \prod_{i=1}^{N} \frac{1}{1+\mu_i C_N}, \qquad P_d^{\text{CA}}(\rho) = \prod_{i=1}^{N} \frac{1}{1+\mu_i C_N/(1+\rho)} \tag{6}$$

This closed form is used to generate theoretical robustness curves versus the number/location of nonhomogeneous cells [4].

2.2.2. GO-CFAR and SO-CFAR

The CFARs split the reference window into two halves of size $n = N/2$ and define

$S_A = \sum_{i=1}^{n} X_i$ and $S_B = \sum_{i=n+1}^{N} X_i$. Under homogeneity, $S_A, S_B \sim \Gamma(n,1)$ are independent. GOCA uses $Y > C_{\mathrm{GO}} \max(S_A, S_B)$ and SOCA uses $Y > C_{\mathrm{SO}} \min(S_A, S_B)$. The corresponding false-alarm probabilities admit finite-sum forms (Erlang shape *n*) [1, 3]:

$$P_{fa}^{\mathrm{SO}}(c) = 2\sum_{m=0}^{n-1} \binom{n-1+m}{m} (2+c)^{-(n+m)}, \qquad P_{fa}^{\mathrm{GO}}(c) = \frac{2}{(1+c)^n} - P_{fa}^{\mathrm{SO}}(c) \tag{7}$$

The calibration constants C_{GO} and C_{SO} are obtained by solving $P_{fa}^{\mathrm{GO}}(C_{\mathrm{GO}}) = P_{fa,0}$ and $P_{fa}^{\mathrm{SO}}(C_{\mathrm{SO}}) = P_{fa,0}$. Under H_1 with SNR ρ, the same expressions apply by the exponential scaling:

$$P_d^{\mathrm{GO}}(\rho) = P_{fa}^{\mathrm{GO}}\left(\frac{C_{\mathrm{GO}}}{1+\rho}\right), \qquad P_d^{\mathrm{SO}}(\rho) = P_{fa}^{\mathrm{SO}}\left(\frac{C_{\mathrm{SO}}}{1+\rho}\right) \tag{8}$$

2.2.3. OS-CFAR and TM-CFAR

Let $X_{(k)}$ denote the *k*-th order statistic of X_i. The OS rule is $Y > T_{\mathrm{OS}} X_{(k)}$. For i.i.d. exponential clutter, the Laplace transform of $X_{(k)}$ yields the closed-form products [25]:

$$P_{fa}^{\mathrm{OS}}(t) = \prod_{i=0}^{k-1} \frac{N-i}{N-i+t}, \qquad P_d^{\mathrm{OS}}(\rho) = \prod_{i=0}^{k-1} \frac{N-i}{N-i+t/(1+\rho)} \tag{9}$$

The multiplier T_{OS} is calibrated by solving $P_{fa}^{\mathrm{OS}}(T_{\mathrm{OS}}) = P_{fa,0}$.

TM-CFAR employs a trimmed mean of the reference samples by discarding the *r* smallest and *s* largest samples and forms the trimmed sum:

$$S_{\mathrm{TM}} = \sum_{i=r+1}^{N-s} X_{(i)}, \qquad Y > C_{\mathrm{TM}} S_{\mathrm{TM}} \tag{10}$$

This approach provides robustness against both strong interference and local clutter variations. Using the independent spacing representation for exponential order statistics, P_{fa} and P_d admit product forms [25, 26]:

$$P_{fa}^{\mathrm{TM}}(c) = \prod_{j=1}^{N} \frac{\lambda_j}{\lambda_j + c w_j}, \qquad P_d^{\mathrm{TM}}(\rho) = \prod_{j=1}^{N} \frac{\lambda_j}{\lambda_j + \frac{c}{1+\rho} w_j} \tag{11}$$

where $\lambda_j = N - j + 1$ and $w_j = \max\big(0, (N-s) - \max(r+1, j) + 1\big)$. The constant C_{TM} is obtained from $P_{fa}^{\mathrm{TM}}(C_{\mathrm{TM}}) = P_{fa,0}$. TM and its variates are often used to improve robustness at clutter transitions.

2.2.4. VI-CFAR switching logic

VI-CFAR adaptively selects among different CFAR schemes based on variability and mean-ratio tests applied to the reference data. Depending on the detected statistical characteristics, the threshold reduces to CA-, GO-, or SO-CFAR with the corresponding scaling factors using the variability index (VI) and mean ratio (MR) computed on each half-window [1]. This switching mechanism allows VI-CFAR to adapt to changing clutter conditions while preserving CFAR behavior under homogeneous environments.

For half-window means $\bar{X}_A, \bar{X}_B$ and sample variances $\hat{\sigma}_A^2, \hat{\sigma}_B^2$, define

$$\mathrm{VI}_A = 1 + \hat{\sigma}_A^2 / \bar{X}_A^2, \quad \mathrm{VI}_B = 1 + \hat{\sigma}_B^2 / \bar{X}_B^2, \quad \mathrm{MR} = \bar{X}_A / \bar{X}_B \tag{12}$$

With thresholds K_{VI} and K_{MR} (set as in the simulations), the switching logic used in this work follows Table 1. The overall (P_{fa}, P_d) of VI-CFAR are evaluated by Monte Carlo due to the data-dependent switching probabilities.

Table 1. *VI-CFAR switching logic.*

Leading variability	Lagging variability	Mean-ratio	Environment discrimination	Adaptive threshold
$\mathrm{VI}_A \le K_{\mathrm{VI}}$	$\mathrm{VI}_B \le K_{\mathrm{VI}}$	$\mathrm{MR} \in [1/K_{\mathrm{MR}}, K_{\mathrm{MR}}]$	Homogeneous	$\tau(X) = C_N(S_{\mathrm{A}} + S_{\mathrm{B}})$
$\mathrm{VI}_A \le K_{\mathrm{VI}}$	$\mathrm{VI}_B \le K_{\mathrm{VI}}$	$\mathrm{MR} \notin [1/K_{\mathrm{MR}}, K_{\mathrm{MR}}]$	Clutter edge	$\tau(X) = C_{\mathrm{GO}} \max(S_{\mathrm{A}}, S_{\mathrm{B}})$
$\mathrm{VI}_A > K_{\mathrm{VI}}$	$\mathrm{VI}_B \le K_{\mathrm{VI}}$	-	Interferences in the leading	$\tau(X) = C_n S_{\mathrm{A}}$
$\mathrm{VI}_A \le K_{\mathrm{VI}}$	$\mathrm{VI}_B > K_{\mathrm{VI}}$	-	Interferences in the lagging	$\tau(X) = C_n S_{\mathrm{B}}$
$\mathrm{VI}_A > K_{\mathrm{VI}}$	$\mathrm{VI}_B > K_{\mathrm{VI}}$	-	Interferences in the both leading and lagging windows	$\tau(X) = C_{\mathrm{SO}} \min(S_{\mathrm{A}}, S_{\mathrm{B}})$

3. VARIANCE-REDUCED SIMULATION OF FALSE-ALARM PROBABILITY

3.1. Conditional-expectation estimator

Direct Monte Carlo estimation of the false-alarm probability,

$$P_{fa} = \Pr\big(Y > \tau(X)\big) \tag{13}$$

becomes impractical when P_{fa} is very small. To address this issue, a variance-reduced estimator based on conditional expectation is employed [19].

Under H_0, the CUT output Y is independent of the reference data X. Conditioning on X, the false-alarm probability can be written as

$$P_{fa} = \mathbb{E}_X\Big[\Pr\big(Y > \tau(X) \mid X\big)\Big] \tag{14}$$

Since Y is exponentially distributed,

$$\Pr\big(Y > \tau(X) \mid X\big) = \exp\left(-\frac{\tau(X)}{\mu_{\mathrm{CUT}}}\right) \tag{15}$$

Thus,

$$P_{fa} = \mathbb{E}_X\left[\exp\left(-\frac{\tau(X)}{\mu_{\mathrm{CUT}}}\right)\right] \tag{16}$$

This estimator eliminates the need to generate CUT samples and significantly reduces variance, enabling reliable estimation of false-alarm probabilities down to 10^{-10} with a moderate number of trials.

3.2. Confidence-interval construction

To quantify estimation uncertainty, confidence intervals are constructed based on the sample mean of

$$w = \exp\left(-\frac{\tau(X)}{\mu_{\mathrm{CUT}}}\right) \tag{17}$$

Assuming a sufficiently large number of trials, a normal approximation is used to obtain 95% confidence intervals. This approach provides a practical measure of numerical stability without excessive computational overhead [22].

3.3. Variance reduction efficiency and computational complexity

To explicitly quantify the advantages of the proposed methodology, we compare the computational efficiency of the conditional-expectation (CE) estimator against the standard Monte Carlo (SMC) approach. For a target false-alarm probability P_{fa}, the SMC estimator requires generating binary detection events, resulting in a theoretical variance of $Var_{SMC} \approx P_{fa} / N_{trails}$. In contrast, the CE estimator evaluates the conditional probability directly, smoothing out the binary fluctuations and yielding a substantially smaller variance for the same number of trials.

Table 2 summarizes the required trials, estimator variance, and the normalized execution time required to achieve a relative error of less than 10% (for a 95% confidence interval) when estimating $P_{fa} = 10^{-4}$ and $P_{fa} = 10^{-6}$ under homogeneous clutter using CA-CFAR ($N = 36$).

Table 2. *Performance comparison between SMC and proposed CE estimators.*

Estimator	Target P_{fa}	Required Trials (N_{trials})	Estimator Variance	Normalized Execution Time
SMC	10^{-4}	1.00×10^{6}	1.00×10^{-10}	100.00% (Baseline)
Proposed CE	10^{-4}	0.56×10^{3}	1.00×10^{-10}	0.06%
SMC	10^{-6}	1.00×10^{8}	1.00×10^{-14}	9953.93%
Proposed CE	10^{-6}	4.41×10^{3}	1.00×10^{-14}	0.45%

As demonstrated, evaluating extremely small P_{fa} via SMC becomes computationally prohibitive due to the immense number of trials required to reach statistical stability (requiring nearly a 100-fold increase in execution time to estimate $P_{fa} = 10^{-6}$). The proposed CE estimator reduces the required trials by several orders of magnitude, accelerating the convergence rate and drastically reducing the computational burden while maintaining rigorous theoretical consistency.

4. SIMULATION SETUP AND NUMERICAL RESULTS

4.1. Parameter settings and clutter-edge modeling

Unless otherwise stated, the parameters used throughout the simulations are given in Table 3.

Table 3. *Parameter settings for the simulations.*

Parameter	Value
Number of reference cells	$N = 36$ ($n = 18$ per side)
Design false-alarm probability	$P_{fa,0} = 10^{-4}$
SNR range	−5 dB to 30 dB
Interference-to-noise ratio (INR)	INR = SNR in multi-target scenarios
Clutter-to-noise ratio (CNR)	CNR = 10 dB and 20 dB in clutter-edge scenarios
The rank parameter of OS-CFAR	$k = 27$ ($3N/4$)
The trim parameters of TM-CFAR	$r = 2$ and $s = 2$
The parameters of VI-CFAR	$K_{VI} = 4.76$ and $K_{MR} = 1.806$ [13, 27]

Regarding the VI-CFAR parameters, the switching thresholds $K_{VI} = 4.76$ and $K_{MR} = 1.806$ are explicitly chosen based on the derivations in [13, 27] for a design false-alarm probability of 10^{-4} and a total reference window size of $N = 36$. These specific values are statistically optimized to ensure that under homogeneous clutter conditions, the VI logic predominantly defaults to the CA-CFAR branch. This structural design enables VI-CFAR to accurately maintain the nominal false-alarm rate, which is corroborated by the baseline validation results presented later in Section 4.2.

Clutter edges are modeled by introducing a step change in the clutter mean across the reference window. Let M_c denote the number of high-clutter cells in the reference window. The reference cell means are defined as

$$\mu_i = \begin{cases} \beta, & i \le M_c \\ 1, & i > M_c \end{cases} \tag{18}$$

where β denotes the clutter-to-noise ratio.

To ensure physical consistency, an edge-aware CUT model is employed:

$$\mu_{\mathrm{CUT}} = \begin{cases} 1, & M_c \le N/2 \\ \beta, & M_c > N/2 \end{cases} \tag{19}$$

This model captures the transition of the CUT between clutter regions as the edge position varies and enables correct interpretation of false-alarm inflation and target masking effects.

4.2. Validation under homogeneous clutter

We first validate the proposed simulation framework under homogeneous clutter (M_I = 0). Unless otherwise stated, the number of reference cells is N = 36 (two-sided with n = 18 per side) and the design false-alarm probability is $P_{fa,0} = 10^{-4}$. Detection performance is evaluated over SNR from −5 dB to 30 dB.

Figure 2 compares simulated and theoretical P_d for CA-, GO-, SO-, OS-, and TM-CFAR. Across the full SNR range, the simulated curves closely match the analytical predictions. VI-CFAR is also included for reference.

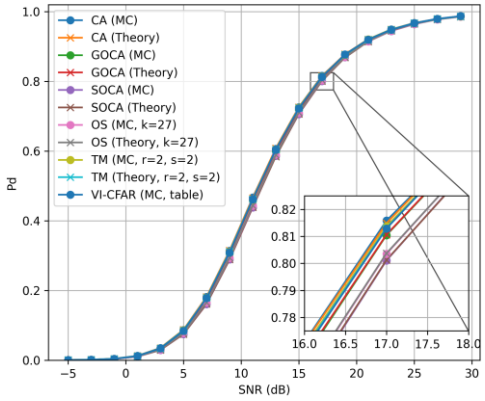

Figure 2. *P_d versus SNR under homogeneous clutter conditions.*

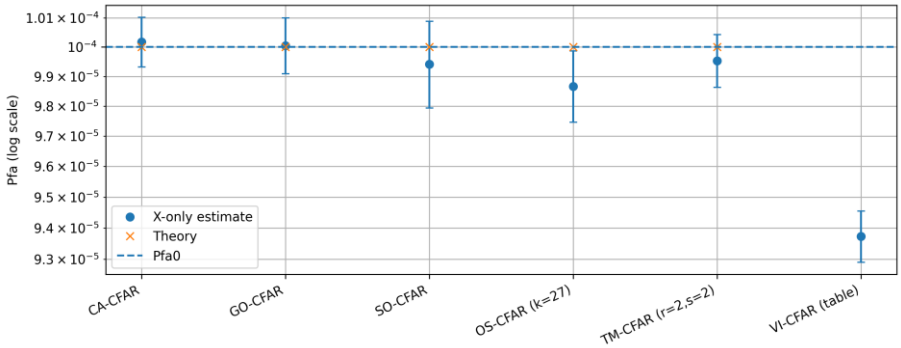

Figure 3. *P_{fa} validation under homogeneous clutter.*

Figure 3 reports the false-alarm validation using the variance-reduced estimator and 95%

confidence intervals. For the CFARs, the estimates are consistent with the design value 10^{-4}, with the theoretical P_{fa} lying within the corresponding confidence intervals. For VI-CFAR, a closed-form P_{fa} is not available; the estimated value remains on the same order as $P_{fa,0}$.

Overall, these results confirm that the proposed implementation and threshold calibration reproduce the classical homogeneous-clutter behavior, providing a reliable baseline for the subsequent non-homogeneous and multi-target experiments.

4.3. Detection performance with multiple interfering targets

This subsection evaluates the detection performance of several CFAR detectors in the presence of multiple interfering targets within the training window. The clutter is assumed homogeneous and exponentially distributed, while M_I interfering targets are introduced into the leading reference window. Each interfering target follows an exponential distribution with mean $1 + \mathrm{INR}$, where the INR is set equal to SNR of the CUT. The CUT always contains the desired target. Detection probability is examined as a function of SNR for $M_I = 1,2,4$ and 6. For reference, the performance of the ideal Neyman-Pearson (NP) detector is also shown.

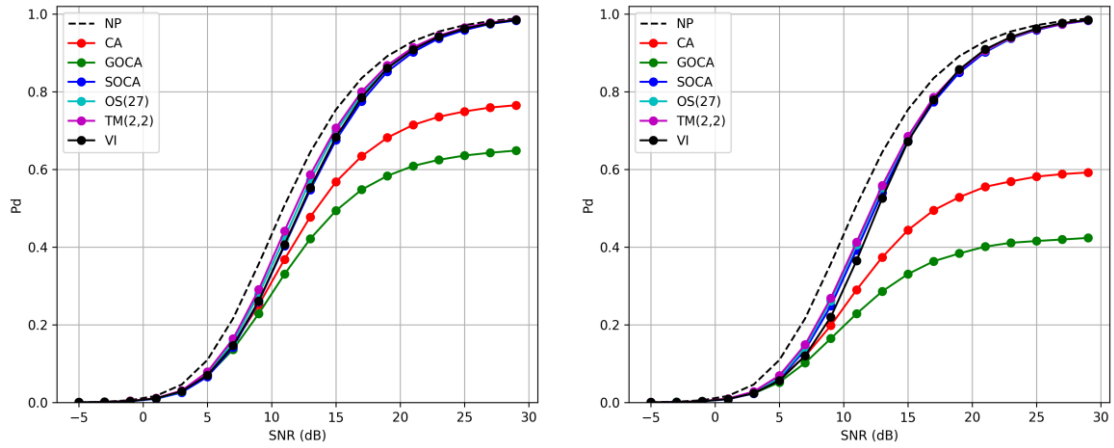

(a) One interfering target. *(b) Two interfering targets.*

Figure 4. *P_d versus SNR with interfering targets in the training window (INR=SNR).*

Figure 4(a) and 4(b) present the detection probability for one and two interfering targets, respectively. Even for small numbers of interferers, CA-CFAR exhibits a noticeable degradation relative to the NP benchmark, while GO-CFAR shows the most pronounced loss among the considered schemes. In contrast, OS-CFAR and TM-CFAR maintain significantly higher detection probabilities, with performance remaining close to that of the NP detector over a wide SNR range. SO-CFAR and VI-CFAR exhibit intermediate behavior.

The impact of stronger interference is illustrated in Figures 5(a) and 5(b) for $M_I = 4$ and 6. In these cases, the detection performance of CA-CFAR and GO-CFAR deteriorates severely, resulting in substantial SNR loss. SO-CFAR provides partial robustness but still suffers notable degradation as the number of interfering targets increases. OS-CFAR consistently achieves near-NP performance, while TM-CFAR preserves good robustness with a moderate loss. VI-CFAR adapts to the interference conditions and maintains detection performance significantly superior to CA- and GO-CFAR across all scenarios.

Overall, the results confirm that the presence of multiple interfering targets can severely degrade the detection performance of conventional CFAR schemes. Robust detectors based on order statistics or trimming exhibit markedly improved resilience, underscoring their suitability for multi-target environments.

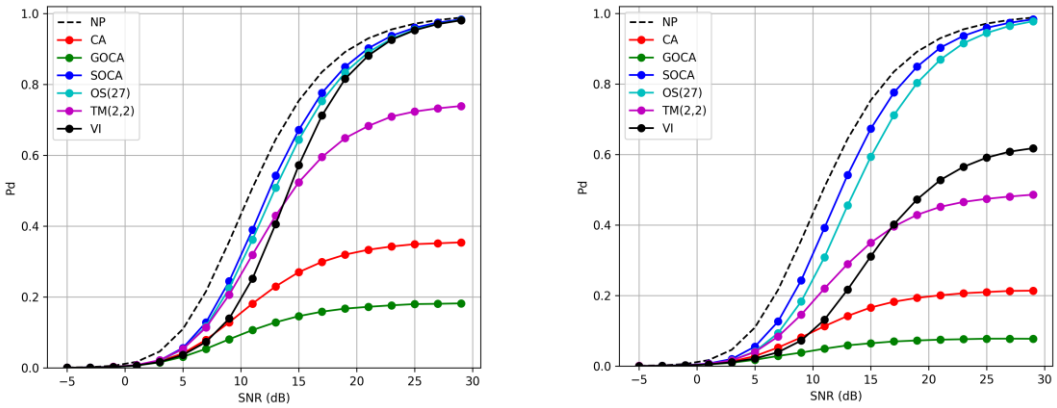

(a) Four interfering targets. *(b) Six interfering targets.*

Figure 5. *P_d versus SNR with interfering targets in the training window (INR=SNR).*

4.4. False-alarm robustness with multiple interfering targets

This subsection evaluates the robustness of several CFAR detectors in terms of false-alarm probability in the presence of multiple interfering targets within the reference window. The CUT contains no target, and the objective is to assess deviations from the design false-alarm probability $P_{fa,0} = 10^{-4}$. All simulations assume homogeneous exponential clutter. Interfering targets are modeled as exponential samples with mean $1 + \mathrm{INR}$, where INR = 15 dB. The number of targets M_I varies from 0 to 6. False-alarm probabilities are estimated using the proposed variance-reduced estimator with 2×10^6 realizations, enabling reliable evaluation down to 10^{-10}.

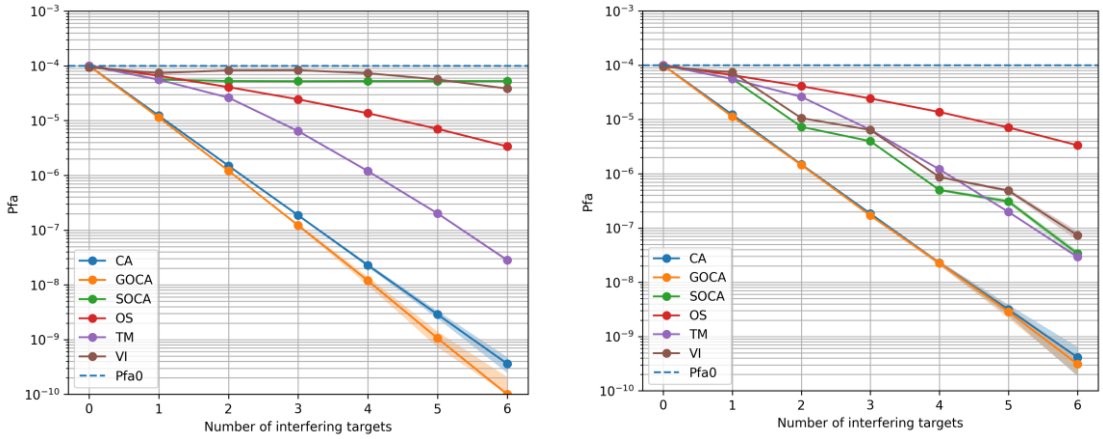

(a) One-sided (all in leading window). *(b) Two-sided (split leading & lagging).*

Figure 6. *Estimated P_{fa} versus number of interfering targets.*

Figure 6(a) shows the false-alarm probability when all interfering targets are located in the leading reference window. While all detectors satisfy the CFAR constraint for $M_I = 0$, significant deviations are observed as M_I increases. CA- and GO-CFAR exhibit a rapid decrease in P_{fa}, indicating a strong loss of the CFAR property. SO-CFAR provides partial robustness but still deviates noticeably for larger M_I.

In contrast, OS- and TM-CFAR maintain substantially more stable false-alarm behavior over the entire range of M_I. VI-CFAR exhibits adaptive behavior, with performance lying between SO- and OS-CFAR depending on the interference level.

The corresponding results for symmetrically distributed interference are shown in Figure 6(b).

Although the degradation becomes more balanced across detectors, CA- and GO-CFAR still experience severe departures from the design false-alarm probability. OS- and TM-CFAR remain the most robust schemes, while VI-CFAR consistently outperforms CA- and GO-CFAR and exhibits robustness comparable to OS-CFAR for moderate interference levels.

These results confirm that conventional CA- and GO-CFAR detectors fail to preserve the CFAR property in multi-target environments, whereas order-statistics-based detectors provide significantly improved robustness. The proposed simulation framework enables accurate characterization of such effects at very low false-alarm probabilities, which is essential for meaningful CFAR performance assessment.

4.5. Effects of clutter edges

This section examines the impact of clutter edges on CFAR performance, with the effects of multiple interfering targets deliberately excluded to isolate the role of clutter non-homogeneity. A step change in clutter power is introduced across the reference window, and an edge-aware CUT model is employed to ensure physical consistency. CNR is set to 10 dB and 20 dB.

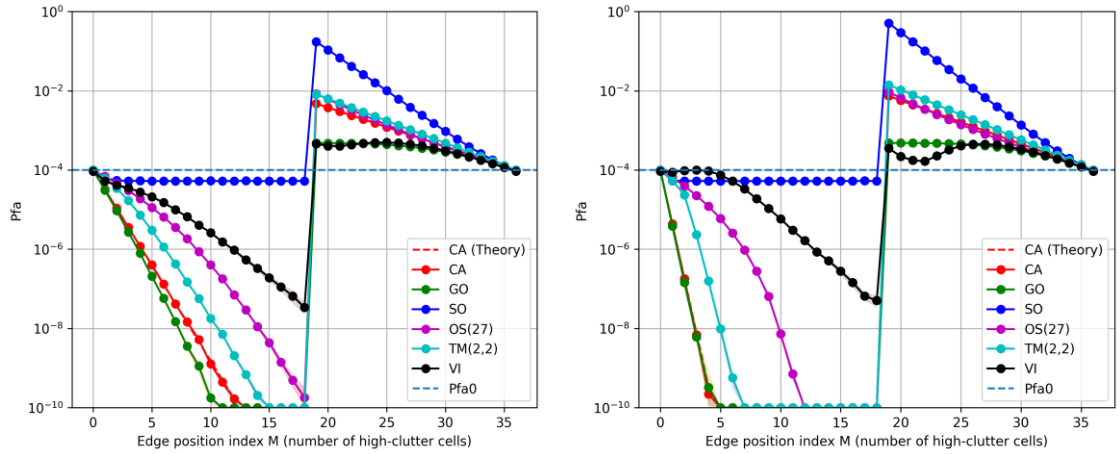

(a) CNR = 10 dB. *(b) CNR = 20 dB.*

Figure 7. *P_{fa} versus clutter edge position.*

Figures 7(a) and 7(b) show the estimated false-alarm probability as a function of the clutter edge position for CNR values of 10 dB and 20 dB, respectively. When the edge is located within the reference window while the CUT remains in the low-clutter region ($M_c \leq N/2$), most detectors exhibit a pronounced reduction in P_{fa}, indicating overly conservative thresholding. This effect is strongest for CA- and GO-CFAR. As the edge moves past the CUT ($M_c \leq N/2$), a sharp increase in false-alarm probability is observed for CA-, GO-, and SO-CFAR, particularly at higher CNR. In contrast, OS-CFAR and TM-CFAR show substantially improved robustness, with P_{fa} remaining much closer to the design value across all edge positions. VI-CFAR mitigates the most severe deviations and provides intermediate robustness between SO- and OS-CFAR. For homogeneous conditions ($M_c = 0$ or $M_c = N$), all detectors recover the design false-alarm probability, confirming consistent behavior under uniform clutter.

To illustrate the impact of clutter edges on detection, the probability of detection is evaluated for three representative configurations: edge after the CUT ($M_c = 12$), edge at the window center ($M_c = 18$), and edge before the CUT ($M_c = 24$), with CNR fixed at 10 dB. The results are shown in Figure 8. When the CUT is located in low clutter while a significant portion of the reference window contains high-clutter cells, CA- and GO-CFAR suffer severe detection loss due to threshold inflation. SO-CFAR provides improved performance in this case, while OS-CFAR and

TM-CFAR maintain relatively stable detection behavior. VI-CFAR adapts to the edge configuration and avoids the extreme target masking observed for conventional CFAR schemes. As the edge moves closer to or past the CUT, detection performance improves for all detectors, although residual losses persist for detectors that are highly sensitive to clutter non-homogeneity.

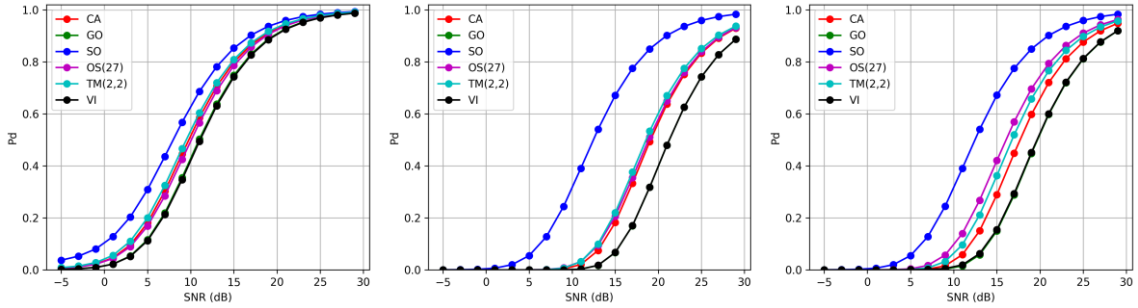

(a) Edge before CUT. *(b) Edge at center.* *(c) Edge after CUT.*

Figure 8. *Detection probability versus SNR with different positions of clutter edge.*

By isolating the effects of clutter edges from those of interfering targets, the present analysis provides a clearer physical interpretation of CFAR behavior in non-homogeneous environments and complements the multi-target performance evaluations presented in Sections 4.3 and 4.4.

5. CONCLUSIONS

This paper developed a theoretically grounded simulation framework for CFAR performance assessment with an emphasis on accuracy, numerical stability, and reproducibility. Correct detector-specific threshold calibration is integrated for CA-, GO-, SO-, OS-, TM-, and VI-CFAR, avoiding common miscalibration pitfalls in conventional studies. For false-alarm evaluation, a variance-reduced estimator enables reliable estimation of extremely small levels without requiring prohibitively large Monte Carlo runs. The framework supports controlled, physics-consistent scenario generation for both homogeneous and non-homogeneous conditions. Overall, the proposed methodology provides a practical and transparent toolset for meaningful CFAR benchmarking and for validating analytical results where available. Future work will extend the framework to more complex statistical environments. Importantly, the mathematical foundation of the conditional-expectation estimator is inherently independent of the specific clutter distribution. As long as the conditional cumulative distribution function of the CUT can be analytically or numerically formulated, the proposed variance-reduction methodology can be readily generalized. Thus, subsequent studies will focus on benchmarking CFAR performance under non-Rayleigh heavy-tailed clutter and Swerling fluctuating target models, leveraging this framework to maintain computational efficiency in the rare-event regime.

REFERENCES

[1]. M. A. Richards, J. A. Scheer, and W. A. Holm. *"Principles of Modern Radar: Basic Principles"*. London, U.K.: IET, (2010). DOI: 10.1049/SBRA021E.

[2]. G. Minkler and J. Minkler. *"CFAR: The Principles of Automatic Radar Detection in Clutter"*. Norwood, MA, USA: Artech House, (1990). DOI: 10.5555/575491.

[3]. H. Rohling. *"Radar CFAR thresholding in clutter and multiple target situations"*. IEEE Trans. Aerosp. Electron. Syst., vol. AES-19, no. 4, pp. 608–621, (1983). DOI: 10.1109/TAES.1983.309350.

[4]. P. P. Gandhi and S. A. Kassam. *"Analysis of CFAR processors in nonhomogeneous background"*. IEEE Trans. Aerosp. Electron. Syst., vol. 24, no. 4, pp. 427–445, (1988). DOI: 10.1109/7.7185.

[5]. G. V. Trunk. *"Range resolution of targets using automatic detectors"*. IEEE Trans. Aerosp. Electron. Syst., vol. AES-14, no. 5, pp. 750–755, (1978). DOI: 10.1109/TAES.1978.308625.

[6]. V. G. Hansen and J. H. Sawyers. *"Detectability loss due to 'greatest of' selection in a cell-averaging CFAR"*. IEEE Trans. Aerosp. Electron. Syst., vol. AES-16, no. 1, pp. 115–118, (1980). DOI: 10.1109/TAES.1980.308885.

[7]. M. E. Shor and N. Levanon. *"Performances of order statistics CFAR"*. IEEE Trans. Aerosp. Electron. Syst., vol. 27, no. 2, pp. 214–224, (1991). DOI: 10.1109/7.78295.

[8]. A. Di Vito, G. Galati, and R. Mura. *"Analysis and comparison of two order statistics CFAR systems"*. IEE Proc. Radar Sonar Navig., vol. 141, no. 2, pp. 109–115, (1994). DOI: 10.1049/ip-rsn:19949886.

[9]. B. Magaz, A. Belouchrani, and M. Hamadouche. *"Automatic threshold selection in OS-CFAR radar detection using information theoretic criteria"*. Prog. Electromagn. Res. B, vol. 30, pp. 157–175, (2011). DOI: 10.2528/PIERB10122502.

[10]. Y. Kong, Y. Wang, J. Cui, and X. Yang. *"Performance prediction of OS-CFAR for generalized Swerling–Chi fluctuating targets"*. IEEE Trans. Aerosp. Electron. Syst., vol. 52, no. 1, pp. 492–500, (2016). DOI: 10.1109/TAES.2015.140967.

[11]. Y. He and H. Meng. *"Performance of a new CFAR detector based on trimmed mean"*. Proc. IEEE Int. Conf. Syst., Man, Cybern. (SMC), (1996). DOI: 10.1109/ICSMC.1996.569880.

[12]. D. Tao, A. P. Doulgeris, and C. Brekke. *"A segmentation-based CFAR detection algorithm using truncated statistics"*. IEEE Trans. Geosci. Remote Sens., vol. 54, no. 5, pp. 2887–2898, (2016). DOI: 10.1109/TGRS.2015.2506822.

[13]. M. E. Smith and P. K. Varshney. *"Intelligent CFAR processor based on data variability"*. IEEE Trans. Aerosp. Electron. Syst., vol. 36, no. 3, pp. 837–847, (2000). DOI: 10.1109/7.869503.

[14]. M. A. Khalighi and M. H. Bastani. *"Adaptive CFAR processor for nonhomogeneous environments"*. IEEE Trans. Aerosp. Electron. Syst., vol. 36, no. 3, pp. 889–897, (2000). DOI: 10.1109/7.869508.

[15]. A. Sarma and D. W. Tufts. *"Robust adaptive threshold for control of false alarms"*. IEEE Signal Process. Lett., vol. 8, no. 9, pp. 261–263, (2001). DOI: 10.1109/97.948451.

[16]. V. Anastassopoulos and G. Lampropoulos. *"Optimal CFAR detection in Weibull clutter"*. IEEE Trans. Aerosp. Electron. Syst., vol. 31, no. 3, pp. 1161–1172, (1995). DOI: 10.1109/7.366292.

[17]. G. V. Weinberg. *"Constant false alarm rate detectors for Pareto clutter models"*. IET Radar Sonar Navig., vol. 7, no. 2, pp. 153–163, (2013). DOI: 10.1049/iet-rsn.2011.0374.

[18]. G. V. Weinberg. *"Formulation of a generalised switching CFAR with application to X-band maritime surveillance radar"*. SpringerPlus, vol. 4, art. no. 574, (2015). DOI: 10.1186/s40064-015-1347-2.

[19]. S. Asmussen and P. W. Glynn. *"Stochastic Simulation: Algorithms and Analysis"*. New York, NY, USA: Springer, (2007). DOI: 10.1007/978-0-387-69033-9.

[20]. V. Elvira and L. Martino. *"Advances in Importance Sampling"*. Wiley StatsRef: Statistics Reference Online, pp. 1–14, (2021). DOI: 10.1002/9781118445112.stat08284.

[21]. N. B. Rached, E. V. Schwerin, G. Shaimerdenova and R. Tempone. *"Importance sampling for rare event tracking within the ensemble Kalman filtering framework"*. Statistics and Computing, vol. 36, no. 1, pp. 1–35, (2025). DOI: 10.1007/s11222-025-10736-1.

[22]. C. J. Clopper and E. S. Pearson. *"The use of confidence or fiducial limits illustrated in the case of the binomial"*. Biometrika, vol. 26, no. 4, pp. 404–413, (Dec. 1934). DOI: 10.1093/biomet/26.4.404.

[23]. J. Xue, S. Xu, and P. Shui. *"Near-optimum coherent CFAR detection of radar targets in compound-Gaussian clutter with inverse Gaussian texture"*. Signal Processing, vol. 166, (2020). DOI: 10.1016/j.sigpro.2019.07.029.

[24]. K. Zebiri et al. *"CFAR detection using two scale invariant functions in heterogeneous Weibull clutter"*. SIViP, vol. 18, pp. 7285–7291, (2024). DOI: 10.1007/s11760-024-03393-w.

[25]. H. A. David and H. N. Nagaraja. *"Order Statistics"*. 3rd ed. Wiley, (2003).

[26]. M. B. El Mashade. *"Detection performance of the trimmed-mean CFAR processor with noncoherent integration"*. IEE Proc. Radar, Sonar and Navigation, vol. 142, no. 1, pp. 18–24, (1995). DOI: 10.1049/ip-rsn:19951626.

[27]. X. Chunbo and L. Yifan. *"Adaptive Constant False Alarm Rate Detector Based on Long Short-term Memory Network"*. Radioengineering, vol. 34, no. 1, pp. 132–142, (2025). DOI: 10.13164/re.2025.0132.

TÓM TẮT

Mô phỏng giảm phương sai để đánh giá chất lượng các bộ phát hiện CFAR trong môi trường không đồng nhất

Bài báo này trình bày một phương pháp mô phỏng chính xác và ổn định để đánh giá chất lượng của các bộ phát hiện có ổn định xác suất báo động lầm (CFAR) trong điều kiện nhiễu không đồng nhất và mục tiêu gây nhiễu. Phương pháp đề xuất kết hợp việc hiệu chuẩn ngưỡng nhất quán về mặt lý thuyết cho các bộ CFAR và trình bày một bộ ước lượng giảm phương sai cho xác suất báo động lầm dựa trên kỳ vọng có điều kiện. Cách tiếp cận này cho phép ước lượng tin cậy các xác suất báo động lầm cực nhỏ mà không tốn quá nhiều chi phí tính toán. Các kịch bản sườn nhiễu được mô hình hóa bằng cách sử dụng mô hình nhận biết sườn, đảm bảo tính nhất quán vật lý giữa cửa sổ tham chiếu và ô kiểm tra. Các kết quả số đã được xác nhận trong điều kiện nhiễu đồng nhất và chứng minh hiệu quả của nó trong việc phân tích độ bền vững của bộ phát hiện khi có nhiều mục tiêu gây nhiễu và các sườn nhiễu. Kết quả này cung cấp cái nhìn sâu sắc về ưu điểm và hạn chế tương đối của các bộ CFAR, đồng thời nhấn mạnh tầm quan trọng của kỹ thuật mô phỏng chính xác để so sánh chất lượng phát hiện có ý nghĩa trong các môi trường ra đa thực tế.

Từ khóa: CFAR; Trung bình-cắt; Chỉ số-biến thiên; Monte Carlo; Giảm phương sai.