An Edge-AI Driven Multi-Layered Air Defense Framework for Real-Time Cost Asymmetric Threat Prioritization

Najmin Akter*
Dept. of Computer Science & Engineering
Jatiya Kabi Kazi Nazrul Islam University
Mymensingh, Bangladesh
najmina757@gmail.com

Younush Sheikh
Dept. of Computer Science & Engineering
Jatiya Kabi Kazi Nazrul Islam University
Mymensingh, Bangladesh
younushsheikh1513@gmail.com

Rahul Raha
Dept. of Computer Science & Engineering
Jatiya Kabi Kazi Nazrul Islam University
Mymensingh, Bangladesh
arishrahul99@gmail.com

Kusuma Akter
Dept. of Computer Science & Engineering
Jatiya Kabi Kazi Nazrul Islam University
Mymensingh, Bangladesh
kusumaakter2003.kulkandi@gmail.com

Md. Sujan Ali
Dept. of Computer Science & Engineering
Jatiya Kabi Kazi Nazrul Islam University
Mymensingh, Bangladesh
sujan_cse@jkkniu.edu.bd

***Abstract*—Unmanned Aerial Vehicles (UAVs) are becoming increasingly common, creating new challenges in maintaining airspace security. As the number of drones in the sky continues to rise, quickly identifying potential threats has become more difficult. Traditional radar and defense systems can be slow to respond and may also mistake harmless objects, such as birds, for potential threats, especially when they are flying at low altitudes. To address these challenges, we propose a smart, multi-layered defense system based on Edge-AI that can uncover and prioritize potential threats in real time. The proposed system uses low-cost hardware, including ESP32-CAMs and ultrasonic sensors, along with a lightweight MobileNetV2 model deployed at the edge. This edge-based design distinguishes between birds and aircraft without relying heavily on cloud-based processing. In addition to object detection, the system introduces a Cost-Asymmetric Threat Prioritization(CATP) algorithm. The algorithm considers several factors, including the target's flight trajectory, potential payload risk, and the feasibility of interception, to calculate a Final Threat Score(FTS). This score helps the system identify the targets that should receive the highest priority. The visual classification model performed the accuracy of 86.73% and an F1-score of 90.04%. The complete system achieved a response time of approximately 145 milliseconds, equivalent to around 7 frames per second, while utilizing only 3.5% of CPU resources. Besides, incorporating the CATP algorithm reduced the simulated interceptor engagement cost by 61.1%. These results suggest that the proposed approach can provide an efficient, affordable and practical solution for real-time UAV threat detection and prioritization.**

***Index Terms*—Edge-AI, MobileNetV2, Cost-Asymmetric Threat Prioritization (CATP), Unmanned Aerial Vehicles (UAVs), Real-Time Air Defense, ESP32-CAM.**

I. Introduction

The rapid growth of low-cost UAVs poses new challenges for air defense: coordinated drone swarms can overwhelm defensive resources because inexpensive UAVs face far costlier interceptors [3], [4], [12], [21]. Traditional decision-support systems rely mainly on kinematic information and classical estimation [1], [2], [16], which struggle in dense swarms where UAVs have similar speeds and distances, making it hard to separate real threats from decoys or birds [6], [8]. Recent work turns to AI-based, multisensor-fused threat assessment [7], [9], [14], with lightweight CNNs and edge computing enabling low-latency, on-node visual inference [4], [5], [17]. Yet perception-accurate systems [15], [17], [18], [23] rarely produce cost-aware engagement decisions, cost/priority-aware frameworks [2], [13], [22] ignore payload severity or destructibility, and comprehensive AI architectures [9] remain conceptual rather than edge-validated.

To address this gap, we propose an Edge-AI-driven, multi-layered air defense framework built around a novel *Cost-Asymmetric Threat Prioritization (CATP)* algorithm. CATP fuses kinematic and CNN-derived features (payload severity, destructibility) with a trajectory tie-breaker and a Cost-Asymmetric Ratio (CAR) to select among three response layers.

The novelty lies not in the MobileNetV2 classifier but in the decision layer behind it. Compared with the studies in Table I, CATP is, to our knowledge, the first framework that (i) combines payload severity, destructibility and trajectory alignment into one score separating targets with nearly identical speed and distance, and (ii) uses the CAR twice, inside the score and as a guard that reserves expensive interceptors for worthwhile targets, so threat and cost are decided together. The key contributions are: (i) a multi-factor threat score resolving the tie-breaking left unsolved by kinematic-only methods; (ii) a tunable CAR linking engagement to economic efficiency; and (iii) a tri-layered response strategy deployable on low-power Edge-AI hardware, validated end-to-end for latency, resource footprint, and cost.

The rest of the paper is organized as follows: Section II reviews relevant literature, Section III describes the proposed CATP algorithm, Section IV presents the system architecture and experimental setup, Section V reports the experimental evaluation and results, and Section VI concludes with future

979-8-3195-4477-3/26/$31.00 ©2026 IEEE

work.

II. Literature Review

Classical air-defense decision-support relies on kinematic feature extraction and control-theoretic tracking [1], [2], [16], effective for isolated targets but unable to resolve threat priority when multiple UAVs share near-identical kinematics [6], [8]. Table I summarizes representative studies against the dimensions relevant to cost-aware, multi-layered engagement.

Most recently, a 2026 study reframed drone/bird discrimination itself as a trajectory-classification problem using a bidirectional LSTM [24], and the 2025 Drone-vs-Bird Grand Challenge at IJCNN benchmarked 16 detection algorithms on this same discrimination task [25]; both studies show that reliable bird filtering is still an open problem. At the architecture level, Iman et al. [3] and Arif [12] compared fielded multi-layered systems (Iron Dome, NASAMS, Hisar-O/A, S-400), confirming the need for tiered, cost-conscious response but offering only descriptive analysis. Hall and Llinas [7] established the basis of sensor fusion, and Shi et al. [4] showed that edge processing cuts latency and energy versus cloud offloading. Zhang et al. [8] applied a Deep Belief Network to post-hoc effectiveness evaluation rather than real-time per-target decisions. Edge-AI-deployable perception is advanced by Howard et al.'s MobileNet [5]; Sharma et al. [10] and Subbulakshmi et al. [11], who prototyped low-cost ultrasonic/ESP32-CAM and multimodal CNN/RNN sensing without reported accuracy or latency figures; and Bayizere [18] and Suwapong et al. [23], whose multimodal fusion reached 96.1% and 95.6% detection accuracy respectively but without an explicit interceptor-cost ratio. Ye et al. [19] confirmed real-time multimodal edge-inference feasibility for localization rather than threat scoring, and Kabir et al. [26] extended dual-stream detection to payload identification with YOLOv11n, again without a cost-aware engagement decision. At the swarm-engagement level, Pandey and Muniraj [13], Hage et al. [22], and Wu et al. [20] proposed criticality-driven assignment, Poisson-prioritized guidance, and PPO-based path planning respectively, while Zhai et al. [21] noted insufficient integration of adversarial threat modeling with resource-aware engagement – none incorporating CNN-derived payload severity or destructibility.

As Table I shows, perception-accurate systems, including the most recent 2025–2026 bird/drone discrimination work [24], [25], rarely reach a cost-aware decision; priority-aware frameworks ignore payload/destructibility; and broad architectures are not tested at the edge. CATP addresses this gap by scoring threat and cost together, as the last column of Table I summarizes.

III. Proposed Cost-Asymmetric Threat Prioritization (CATP) Algorithm

The proposed framework integrates multi-sensor perception, sensor fusion, edge-based AI inference and an autonomous CATP engine that orchestrates a tri-layered response.

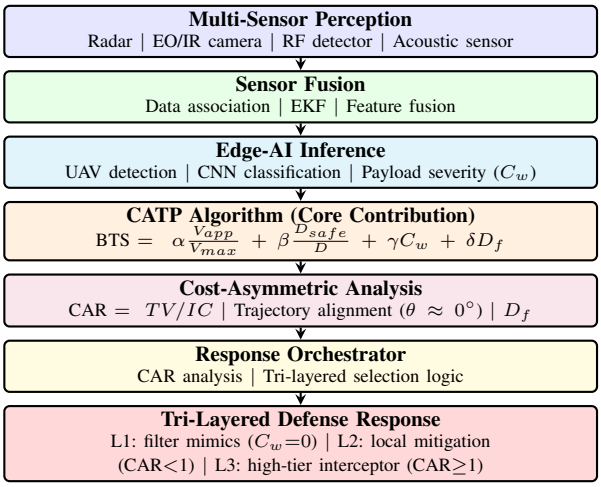


Fig. 1: Proposed seven-layer Edge-AI driven air defense architecture, aligned with the CATP algorithm.

Information moves from top to bottom: sensing and fusion feed the edge classifier, CATP turns its output into scores, and the orchestrator chooses one of the three responses. Each layer is described in Section III.

Fig. 1 presents the seven-layer architecture: biological mimics are filtered when $C_w = 0.0$ (Layer 1), low-cost swarms are neutralized locally when $\mathrm{CAR} < 1.0$ (Layer 2), and critical threats are escalated to high-tier interceptors when $\mathrm{CAR} \geq 1.0$ (Layer 3).

A. Multi-Sensor Perception and Fusion

The perception layer aggregates heterogeneous sensor data for UAV detection and classification, supplying the spatial and visual features for BTS computation: radar provides approach velocity (V_{app}) up to 10 km range; the EO/IR camera enables CNN-derived payload-severity classification (C_w) and bird/UAV discrimination; RF sensing identifies command-and-control links for intent estimation; and an acoustic sensor detects quiet UAVs within 1 km for perimeter filtering. In the fusion layer, data association correlates detections across sensors using Mahalanobis distance; an Extended Kalman Filter (EKF) performs trajectory estimation, enabling monocular trajectory alignment ($\theta \approx 0^\circ$ for direct collision-course detection); and feature fusion combines radar-derived kinematics (V_{app}, D) with camera-derived visual features (C_w) into a unified target vector.

B. CATP Algorithm (Core Contribution)

The CATP algorithm computes normalized threat scores by integrating spatial, visual, trajectory, payload, and destructibility features. Its key components are:

1) **Base Threat Score (BTS):** fuses spatial features (V_{app}, D) and visual features (C_w, D_f) into a normalized score $\mathrm{BTS} \in [0, 1]$:

$$\mathrm{BTS} = \alpha\left(\frac{V_{app}}{V_{max}}\right) + \beta\left(\frac{D_{safe}}{D}\right) + \gamma(C_w) + \delta(D_f) \quad (1)$$

where $\alpha + \beta + \gamma + \delta = 1$ are tunable hyperparameters.

TABLE I: Comparison of Existing Studies and Identified Research Gaps

Study	Perception / Fusion	Threat Priority	Payload / Destructibility	Cost Awareness	Edge Validation	Main Gap
Alqudsi & El-Bayoumi [1]	Kinematic	×	×	×	×	Tracking rather than multi-target prioritization
Tuncer & Cirpan [2]	Kinematic	✓	×	Limited	×	No visual/payload-based assessment
Taşdemir & Atan [6]	ANN + kinematic	✓	×	×	×	No multimodal or economic dimension
Arapoglou et al. [14]	Multisensor + AI	Limited	×	×	×	Sensor selection, not engagement priority
Dabbabi & Delleji [15]	Multimodal + GNN	×	×	×	✓	Tracking without engagement decision
Wang et al. [17]	Visual + Edge-AI	×	×	×	✓	Detection without threat prioritization
Hage et al. [22]	Radar + kinematic	✓	×	Limited	Limited	No payload/destructibility assessment
Pandey & Muniraj [13]	Multisensor	✓	×	✓	×	No payload-aware Edge-AI decision
Volkov & Chornobay [9]	Multimodal AI	✓	Limited	Limited	Limited	Broad architecture; limited edge validation
Proposed CATP	**Multimodal + CNN**	✓	✓	✓	✓*	**Unified threat, cost and response prioritization**

✓: fully addressed; Limited: partially addressed; ×: not addressed. *The sensing node is the ESP32-CAM/Arduino Uno prototype, while latency and CPU/RAM were profiled on the development runtime (Section V). The table makes clear that none of the earlier studies covers threat priority, payload/destructibility and cost together, and this is the gap that CATP fills.

2) **Trajectory Alignment (θ):** the collision-course angle from monocular vision; $\theta \approx 0^\circ$ indicates a direct head-on approach and acts as a tie-breaker among kinematically identical threats.
3) **Payload Severity (C_w):** a CNN-derived confidence quantifying potential damage severity ($C_w \in [0,1]$; e.g. bird $= 0.0$, commercial drone ≈ 0.5, military-grade platform ≈ 1.0).
4) **Destructibility (D_f):** a learned estimate of target vulnerability, $D_f \in [0,1]$; lower values favor low-cost mitigation, higher values require high-tier interceptors.

CATP couples BTS with a Cost-Asymmetric Ratio (CAR) to select the most economically viable engagement:

$$\mathrm{CAR} = \frac{TV}{IC} \tag{2}$$

where TV is the estimated target value and IC is the high-tier interceptor cost. Algorithm 1 formalizes the per-frame decision: once a target enters the safe perimeter ($D < D_{safe}$), biological mimics are filtered (Layer 1), targets with $\mathrm{BTS} > \tau_{critical}$ are routed by the CAR to Layer 2 or Layer 3, and targets exceeding only τ_{alert} are tracked.

C. Final Threat Score and Threat-Level Mapping

All accepted targets are ranked by a Final Threat Score that adds a trajectory tie-breaker and cost-awareness to BTS:

$$\mathrm{FTS} = \alpha\,\mathrm{BTS} + \beta(1-\theta_n) + \gamma C_w + \delta D_f + \lambda\,\mathrm{CAR} \tag{3}$$

where θ_n is the normalized trajectory alignment and $\alpha, \beta, \gamma, \delta, \lambda$ are tunable weights; the raw score is min–max normalized so that $\mathrm{FTS} \in [0,1]$. Note that C_w and D_f contribute to FTS both indirectly (through BTS) and directly; this intentional reinforcement amplifies payload severity and destructibility as tie-breaking factors during final prioritization, and the $[0,1]$ range is guaranteed by the final normalization rather than by the raw coefficient sum.

The normalized FTS is mapped to the three response layers used consistently throughout this paper (TABLE II).

After filtering, targets are ranked by FTS and the response coordinator assigns each one the most cost-efficient response of TABLE II.

Algorithm 1 Cost-Asymmetric Threat Prioritization (CATP)

Require: Ultrasonic distance D, camera frame F, thresholds $\tau_{alert}, \tau_{critical}$, safe distance D_{safe}, max. velocity V_{max}, interceptor cost IC
Ensure: Engagement decision ED (L1: Filter / L2: Local Mitigation / L3: High-Tier Escalation)

while true **do**
  if $D < D_{safe}$ **then**
    $V_{app} \leftarrow |D_t - D_{t-1}|/\Delta t$
    $(Class, C_w, TV, D_f) \leftarrow \mathrm{EdgeAI_Inference}(F)$
    if $C_w = 0.0$ **then**
      continue {L1: filter biological mimic}
    end if
    Compute BTS via Eq. (1); $\mathrm{CAR} \leftarrow TV/IC$
    if $\mathrm{BTS} > \tau_{critical}$ **then**
      if $\mathrm{CAR} < 1.0$ **then**
        $ED \leftarrow$ L2: Local Edge Mitigation; activate turret/alarm
      else
        $ED \leftarrow$ L3: Escalate to High-Tier Defense
      end if
    else if $\mathrm{BTS} > \tau_{alert}$ **then**
      $ED \leftarrow$ Track and Monitor
    end if
  end if
end while

TABLE II: Threat-Level Mapping and Tri-Layered Response Strategy

FTS Range	Level	Response	CAR
$\mathrm{FTS} \geq 0.80$	High	L3: High-tier interceptor	≥ 1.0
$0.50 \leq \mathrm{FTS} < 0.80$	Medium	L2: Edge mitigation (jam/turret)	< 1.0
$\mathrm{FTS} < 0.50$	Low	L1: Passive track / filter	–

IV. System Architecture and Experimental Setup

The proposed decentralized edge-computing architecture optimizes processing at the tactical edge by integrating hardware sensing, visual computing, and tiered actuation.

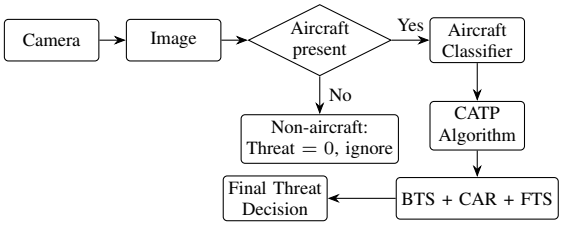


Fig. 2: Aircraft detection and threat assessment pipeline.

As shown in Fig. 2, each camera frame is first checked for an aircraft; frames without one receive zero threat, while confirmed aircraft pass through the aircraft classifier and the CATP algorithm, where BTS, CAR and FTS are combined into the final threat decision.

A. Hardware, Kinematics, and Dataset

The physical perimeter node uses a low-power, heterogeneous sensor stack: an HC-SR04 ultrasonic transducer monitors a 400 cm perimeter via acoustic echo returns; an ESP32-CAM module streams live video frames (F) to a local control station; and an Arduino Uno performs high-frequency sensor polling and signal filtering before forwarding compiled features to the edge node for AI classification. Distance is computed via Time-of-Flight, $D = (T \times v_s)/2$, where T is echo propagation time and $v_s = 343$ m/s; approach velocity uses back-difference differentiation, $V_{app} = |D_t - D_{t-1}|/\Delta t$.

A custom visual dataset of 2,750 images (950 bird, 1,800 aircraft/UAV) was partitioned into training (64%), validation (16%), and test (20%) sets using stratified splitting with a fixed random seed (42), so the split is exactly reproducible. A synthetic dataset of 100 heterogeneous aerial targets, spanning plausible kinematic profiles, payload weights, destructibility, and cost values, was generated to evaluate the CATP decision layer in isolation; each attribute is drawn independently and uniformly (distance 0.5–10 km, speed 60–350 km/h, heading 0–180°, payload 0.5–20 kg, destructibility 0.2–1, target value 10–500 k USD, interceptor cost 5–120 k USD), and seed 42 reproduces Table IV exactly. The complete pipeline (MobileNetV2 inference + CATP) was profiled end-to-end for latency and CPU/RAM utilization and compared against a conventional kinematic-only baseline; Table III lists the exact settings behind every reported result.

V. Experimental Evaluation and Results

A. Visual Classifier Performance

The MobileNetV2 module was evaluated on the held-out test set, achieving 86.73% overall accuracy, 88.47% precision, 91.67% recall, and a 90.04% F1-score on the aircraft class. Fig. 3 shows that training accuracy and loss improve steadily, although validation accuracy fluctuates across epochs before settling near its best value, and Fig. 4 confirms strong binary discrimination, correctly identifying 330 aircraft and 147 bird targets – minimizing wildlife-induced false alarms during perimeter surveillance.

TABLE III: Implementation Settings for the Classifier and the CATP Scores

Item	Setting
Backbone	MobileNetV2 (ImageNet weights), ≈2.26 M parameters
Input preprocessing	224×224 RGB, pixel values scaled by MobileNetV2's own preprocessing function
Classifier head	global average pooling, dropout 0.3, single sigmoid unit
Fine-tuning	last 30 backbone layers trainable, remaining layers frozen
Augmentation	random horizontal flip, rotation (factor 0.1), zoom (factor 0.1)
Training	Adam optimizer, learning rate 10^{-4}, batch size 8, 10 epochs, binary cross-entropy loss
Model selection	checkpoint saved on best validation accuracy; early stopping on validation loss (patience 3)
CATP weights	α=0.35, β=0.30, γ=0.20, δ=0.15, λ=0.20
CATP constants	D_{safe}=10 km; V_{max} = largest speed in the target set; level thresholds 0.50 and 0.80
Random seed	42 (image split and synthetic target generation)

B. CATP Simulation and Threat Ranking

A 100-target synthetic simulation was used to assess the ability of CATP to resolve threat-ranking uncertainty in dense, multi-target scenarios. The results in Table IV show that, by considering trajectory alignment and other contextual factors, CATP assigns higher priority to targets that look less threatening by distance or velocity alone but carry significant payload or less favorable interception cost ratios.

TABLE IV: Selected CATP Threat Ranking Results (seed-42 synthetic raid of 100 targets)

Target ID	Dist. (km)	Payload (kg)	Destr.	Cost	BTS	CAR	θ_n	FTS	Threat Level
UAV-060	0.814	14.31	0.581	31337.92	0.652	0.144	0.015	1.000	High
UAV-019	2.429	19.19	0.952	15104.83	0.236	0.066	0.041	0.905	High
UAV-032	1.660	15.74	0.800	28575.21	0.293	0.239	0.094	0.877	High
UAV-043	7.078	16.82	0.800	32104.55	0.067	0.183	0.120	0.748	Medium
UAV-045	2.819	16.85	0.668	109404.61	0.188	0.001	0.092	0.741	Medium
UAV-028	2.575	12.25	0.822	90787.07	0.167	0.000	0.058	0.710	Medium
UAV-002	9.731	14.77	0.285	86759.54	0.045	0.055	0.250	0.497	Low
UAV-098	9.831	18.35	0.630	31145.36	0.019	0.107	0.548	0.496	Low
UAV-047	2.578	4.84	0.806	72655.27	0.169	0.034	0.311	0.496	Low
UAV-022	7.100	19.16	0.935	18420.05	0.066	0.040	0.744	0.494	Low

The table lists ten of the 100 simulated targets, covering all three threat levels of Table II. Destr. is the destructibility D_f, Cost is the high-tier interceptor cost IC in USD, CAR is shown after min–max rescaling to $[0, 1]$, and θ_n is the normalized trajectory alignment.

C. Threat Score Distribution and Operational Efficiency

Fig. 5 illustrates the distribution of FTS values across the dataset, with the targets classified into three response levels according to TABLE II: Low ($\text{FTS} < 0.50$), Medium ($0.50 \leq \text{FTS} < 0.80$), and High ($\text{FTS} \geq 0.80$). This narrows high-level interception to a small group of high-risk targets, helping prevent resource saturation during swarm operations.

Unlike conventional approaches based on distance and speed, CATP also accounts for economic and contextual factors. Fig. 6 makes this concrete on the same seed-42 targets already listed in Table IV: several targets with a low kinematic-only score (speed and proximity alone) get a much higher FTS once payload, destructibility and trajectory alignment are added. For example, UAV-019 rises from 0.203 to an FTS of 0.905 because it carries a heavy payload and flies almost

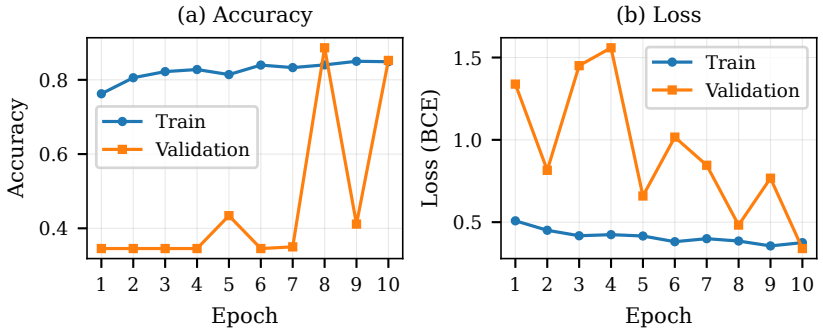


Fig. 3: Training and validation accuracy/loss curves over 10 epochs.

Panel (a) shows accuracy and panel (b) shows binary cross-entropy loss for the training and validation sets over the ten epochs. Training improves steadily, while validation accuracy varies from epoch to epoch before reaching its best value.

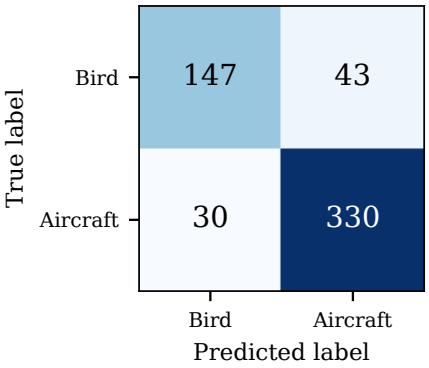


Fig. 4: Birds vs. aircraft confusion matrix on the 550-image held-out test set.

Rows give the true class and columns the predicted class. Out of 190 bird images, 147 are recognized correctly and 43 are taken for aircraft; out of 360 aircraft images, 330 are recognized correctly and 30 are taken for birds.

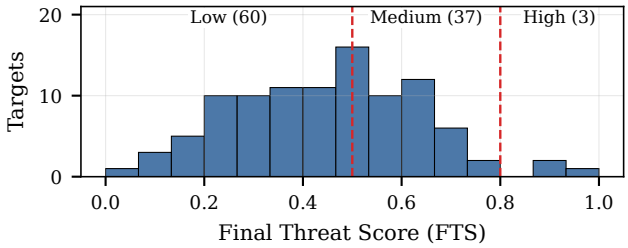


Fig. 5: Distribution of Final Threat Scores generated by CATP.

The histogram covers all 100 targets of the seed-42 raid (Table IV). The dashed lines mark the Low/Medium and Medium/High thresholds of Table II; most targets fall in the Low and Medium bands, and only three reach the High band.

head-on. TABLE V summarizes this qualitative improvement in system capability.

TABLE V: Feature Comparison: Conventional vs. Proposed CATP

Feature / Metric	Conventional	Proposed CATP
Distance & Velocity	✓	✓
Trajectory Intent	×	✓
Payload Severity	×	✓
Destructibility Analysis	×	✓
Cost-Aware Evaluation	×	✓
Multi-Target Tie-Breaking	Limited	Effective

The table summarizes the inputs that each approach takes into account. The measured effect of the additional inputs is reported in Fig. 6 and Table VI.

D. Simulated Cost and Risk-Reduction Comparison

To quantify the operational effect of the tri-layered response, we simulated repeated raids under a fixed, stated engagement model and compared the conventional kinematic-only doctrine (rank by speed and proximity, engage every accepted target with a high-tier interceptor) against the full CATP pipeline (rank by FTS, apply the tri-layered response of TABLE II, and use the CAR guard so that a high-tier interceptor is only spent when $\mathrm{CAR} \geq 1$). Each of 500 independent raids draws a fresh 100-target set with the same distributions as Section IVA; a target's true expected damage is modeled as $p_{hit} \cdot TV \cdot (\text{payload}/20) \cdot s$, where $p_{hit} = \cos^2\theta$ for a heading $\theta < 90°$ (and 0 for a target moving away) and s scales with speed; a high-tier interceptor succeeds with probability 0.95 and a local-mitigation engagement succeeds with probability $0.9(1-D_f)+0.05$, both discounted when the time to impact is very short. We state these assumptions openly and apply them in the same way to both policies, so that the difference comes only from CATP's ranking and tier selection.

TABLE VI: Simulated Damage Neutralized and Engagement Cost, Conventional vs. CATP (mean over 500 raids, 95% confidence interval)

Policy	Damage neutralized	Engagement cost
Conventional (kinematic-only)	$7.7\% \pm 0.4\%$	\$625k ± \$10k
Proposed CATP	$36.3\% \pm 1.1\%$	\$244k ± \$14k

Both differences are statistically significant (paired Wilcoxon signed-rank test, $p < 10^{-80}$, $n = 500$ raids). On average, CATP neutralizes 4.7× more expected damage and spends 61.1% less than the conventional policy.

On the development platform, the complete pipeline (MobileNetV2 + CATP) sustains 145.33 ms end-to-end latency (6.88 FPS) at only 3.5% CPU and 14.3% RAM utilization, and the CATP decision layer itself takes under 8 ms of that total, a cost that suits continuous edge operation. These figures were measured on the Google Colab runtime used for development, not on the ESP32-CAM/Arduino Uno node; profiling on the target hardware is left for future work (Section VI).

E. Discussion

Two results explain the advantage of CATP over the conventional doctrine, beyond a simple reordering of the same decisions. First, Fig. 6 shows that FTS moves heavy-payload, head-on targets that a kinematic-only rule ranks low to the top of the queue. Second, Table VI shows that most of the saving does not come from cheaper interceptors. It comes from not firing a high-tier interceptor at threats that the CAR guard and the Medium/Low tiers send to local mitigation or tracking. This is the behaviour the tri-layered response of Section III was designed for, and the difference is statistically significant. Deciding threat and cost together therefore changes

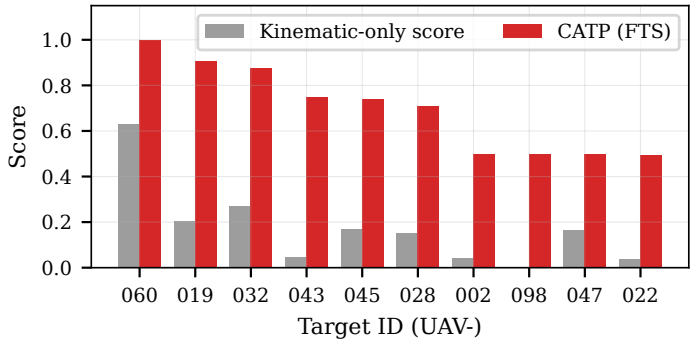


Fig. 6: Kinematic-only score versus CATP's Final Threat Score for the ten targets of Table IV.

Grey bars use speed and proximity only (the V_{app}/V_{max} and D_{safe}/D terms of BTS), and red bars show the full FTS. The gap is largest for targets such as UAV-019 and UAV-043, whose payload and heading matter more than speed and distance.

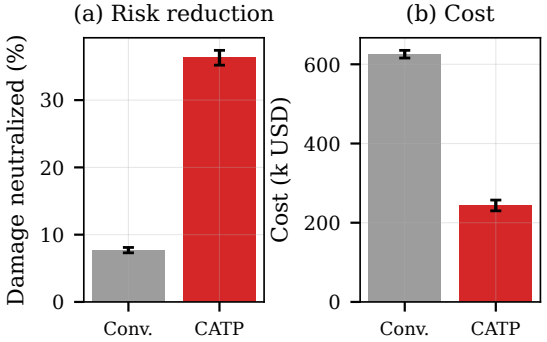


Fig. 7: Simulated damage neutralized and engagement cost, conventional versus CATP.

Each bar is the mean over 500 simulated raids (Conv.: conventional kinematic-only policy), and the error bars give the 95% confidence interval. CATP neutralizes more damage (a) at lower cost (b).

engagement decisions in a way a single kinematic ranking cannot, in line with the gap seen in prior work (Table I).

VI. Conclusion and Future Scope

This paper presented an Edge-AI-driven Multi-Layered Air Defense Framework centered on the proposed Cost-Asymmetric Threat Prioritization (CATP) algorithm for efficient and cost-aware management of UAVs swarm threats. By integrating kinematic information with payload severity, destructibility and trajectory alignment, it enables more informed threat prioritization. The lightweight MobileNetV2-based perception module achieved 86.73% accuracy, while the complete end-to-end pipeline operated at 145.33 ms (6.88 FPS) with only 3.5% CPU and 14.3% RAM utilization. In a 500-raid simulation with a stated engagement model, CATP neutralized 36.3% of the expected damage and spent 61.1% less than the conventional kinematic-only baseline, whose neutralized share was only 7.7%. This shows the potential of CATP as a lightweight, real-time and cost-aware decision aid for Counter-UAS operations. Since the comparison uses synthetic targets and assumed engagement parameters, it demonstrates the working mechanism rather than a field-validated figure. Future work will extend the framework to multi-node and federated Edge-AI deployments, validate it against physical UAV swarms outdoors, develop end-to-end CNN-based estimation of payload severity and destructibility, and examine robustness under sensor noise, communication delays and deceptive UAV behaviors.

References

[1] Y. Sh. Alqudsi and G. M. El-Bayoumi, "Modeling of Target Tracking System for Homing Missiles and Air Defense Systems," *INCAS Bulletin*, 2018.

[2] O. Tuncer and H. A. Cirpan, "Target priority based optimisation of radar resources for networked air defence systems," *IET Radar, Sonar & Navigation*, 2022.

[3] K. F. Iman, R. H. Triharjanto, H. B. Wibowo, and Y. Ruyat, "Comparative Analysis Of A Multi-Layered Weapon System For City Air Defense In The Modern Warfare," *IJHESS*, 2023.

[4] W. Shi, J. Cao, Q. Zhang, Y. Li, and L. Xu, "Edge Computing: Vision and Challenges," *IEEE Internet of Things Journal*, vol. 3, no. 5, pp. 637–646, 2016.

[5] A. G. Howard *et al.*, "MobileNets: Efficient Convolutional Neural Networks for Mobile Vision Applications," *arXiv:1704.04861*, 2017.

[6] S. Taşdemir and M. Atan, "Threat Assessment Model in Air Defense Systems Using Artificial Neural Networks," *Defense and Security Studies*, vol. 7, no. 1, pp. 22–42, 2026.

[7] D. L. Hall and J. Llinas, "An Introduction to Multisensor Data Fusion," *Proceedings of the IEEE*, vol. 85, no. 1, pp. 6–23, 1997.

[8] P. Zhang, K. Feng, J. Gong, X. Yang, and J. Shen, "A Systematic Effectiveness Evaluation Method for Air Defense Operations Based on Deep Confidence Networks," *Research Square (preprint)*, 2024.

[9] A. Volkov and V. Chornobay, "Integrated Algorithmic Model of Air Defense Systems Based on Artificial Intelligence," in *Proc. V All-Ukrainian Sci.-Tech. Conf. Computer Games and Multimedia*, 2025.

[10] K. J. S. Sharma, Y. J. S. Babu, S. T. Saran, and S. Geerthana, "Automatic Smart Defense Based Surveillance Security System Using IoT," *J. Electrical Engineering and Automation*, vol. 7, no. 2, pp. 124–137, 2025.

[11] A. Subbulakshmi, M. Sandhiya, S. Palanivel, D. Denceli, B. S. Shivani, and M. Chezhiyan, "Deep Learning-Based Classification for Border Defence Mechanism," 2024.

[12] S. Arif, "India's Acquisition of the S-400 Air Defense System," *Journal of Indo-Pacific Affairs*, Fall 2021.

[13] S. Pandey and D. Muniraj, "Critical Infrastructure Defense Against Aerial Swarms Under Sensing Uncertainty: Online Allocation With Finite-Time Guarantees," *arXiv preprint (eess.SY)*, 2026.

[14] F. Arapoglou, P. Zacharia, and M. Papoutsidakis, "Intelligent Counter-UAV Threat Detection Using Hierarchical Fuzzy Decision-Making and Sensor Fusion," *Sensors (MDPI)*, 2025.

[15] K. Dabbabi and T. Delleji, "Graph neural network-tracker: a graph neural network-based multi-sensor fusion framework for robust unmanned aerial vehicle tracking," *Visual Computing for Industry, Biomedicine, and Art*, 2025.

[16] N. Zhu, F. Zhong, X. Lei, G. Niu, H. Xie, and Y. Zhang, "Situation Awareness and Tracking Algorithm for Countering Low-Altitude Swarm Target Threats," *Remote Sensing*, vol. 17, no. 7, 2025.

[17] H. Wang, Y. Qu, Z. Dang, D. Wu, M. Cui, H. Shi, and J. Zhao, "A Ground-Based Visual System for UAV Detection and Altitude Measurement: Deployment and Evaluation of Ghost-YOLOv11n on Edge Devices," *Sensors (MDPI)*, vol. 26, no. 1, 2025.

[18] M. Bayizere, "DroneShield-AI: A Multi-Modal Sensor Fusion Framework for Real-Time Autonomous Drone Threat Detection, Behavioral Intent Classification, and Swarm Intelligence in Contested Airspace," *arXiv preprint (cs.CV)*, 2026.

[19] Z. Ye, Y. He, G. Yu, and P. Loskot, "Edge-Assisted Multimodal UAV Localization with Resource-Efficient Compression and Robust Fusion," *arXiv preprint (eess.SP)*, 2026.

[20] Y.-T. Wu, W.-H. Wu, S.-T. Lin, and H.-T. Chiao, "Battlefield Applications of Drone Swarms Based on Proximal Policy Optimization Reinforcement Learning," in *Proc. IEEE ICCBE*, 2025.

[21] L. Zhai, H. Wu, L. Lai, and Z. Gao, "Intelligent Optimization Algorithms for Multi-UAV Path Planning: A Comprehensive Review," *IEEE Access*, vol. 13, 2025.

[22] G. Hage, A. Santoso, and M. Sahal, "Integrated Radar and Missile System with Poisson-Prioritized Threat Management and PPN Guidance for Countering Multiple UAV Threats," *Journal of Robotics and Control (JRC)*, 2025.

[23] C. Suwapong, C. Wechtaisong, and P. Pao-la-or, "Air Threat Defense System from Unidentified Drones Using Multi-Sensor AI," in *Proc. 15th Int. Conf. IEEA*, Osaka, Japan, 2026.

[24] A. Nesteruk, V. Nikitin, Y. Albrekht, Ł. Ścisło, D. Grela, and P. Król, "Distinguishing a Drone from Birds Based on Trajectory Movement and Deep Learning," *Sensors*, vol. 26, no. 3, p. 755, 2026.

[25] A. Coluccia *et al.*, "The Drone-vs-Bird Detection Grand Challenge at IJCNN 2025," in *Proc. Int. Joint Conf. Neural Networks (IJCNN)*, 2025, pp. 1–8.

[26] S. Kabir *et al.*, "SpectraSentinel: LightWeight Dual-Stream Real-Time Drone Detection, Tracking and Payload Identification," *arXiv:2507.22650*, 2025.