Research Square

Preprints are preliminary reports that have not undergone peer review.
They should not be considered conclusive, used to inform clinical practice,
or referenced by the media as validated information.

Uncertainty-Aware Energy-Optimized Adaptive Swarm Coordination for AI-Enabled Multi-Surveillance

NAZMUL HASAN ANIK CHAWDHUY
Aviation and Aerospace University
Mahabub Sultan
mahabubsultan2@gmail.com

SUNY Buffalo State University

Research Article

Keywords: multi-UAV surveillance, swarm coordination, belief map, energy-aware planning, return-to-base feasibility, uncertainty quantification, receding-horizon optimization

Posted Date: May 21st, 2026

DOI: https://doi.org/10.21203/rs.3.rs-9239305/v1

License: This work is licensed under a Creative Commons Attribution 4.0 International License.
Read Full License

Additional Declarations: No competing interests reported.

Uncertainty-Aware Energy-Optimized Adaptive Swarm Coordination for AI-Enabled Multi-Surveillance

Nazmul Hasan Anik Chawdhury [1*], *Mahabub Sultan* [2]

[1*]EMAIL: anikchawdhury@aaub.edu.bd ORCID ID: 0009-0001-7153-0186 Address: Aviation and Aerospace University Bangladesh, Lalmonirhat Sadar, Lalmopnirhar, Bangladesh. Zip Code: 5500

[2]Email: mahabubsultan2@gmail.com Orcid ID: 0009-0004-8589-3270 City: Buffalo, NY State: New York Zip Code: 14212 Short version: Buffalo, NY 14212 Mobile- 412-218-9731 House address: 149 Milburn St, Buffalo, NY 14212

Abstract

Persistent multi-UAV surveillance requires simultaneous management of area coverage, collision avoidance, battery reserve, uncertain detections, and mission continuity. Existing work remains fragmented across fixed-rule swarm coordination, energy-aware planning without calibrated perception, and vision modules that are detached from the control loop. This paper presents uncertainty-aware energy-optimized adaptive swarm coordination (UEO-ASC), a unified framework that couples calibrated threat belief, adaptive interaction fields, local receding-horizon optimization, and prediction-based return-to-base feasibility in a tractable surveillance architecture. The model combines a layered 6-DOF/PID flight stack, Bayesian belief-map updates, state-of-charge propagation, distributed task selection, and energy-feasible local decision making. A reference simulation campaign was conducted on a 1.2 km x 1.2 km surveillance map with six UAVs, four operating scenarios, and four independent seeds per scenario (48 runs total). Against fixed-Boids coordination, UEO-ASC reduced mission energy consumption by 7.2-8.1%, lowered false dispatches by 55.8-92.6%, and improved incident-map F1 score by 47.7-63.7%, while keeping gross coverage within 0.9-2.7 percentage points of the best baseline across scenarios. Relative to an energy-aware rule baseline, UEO-ASC also maintained lower false dispatch rates and higher F1 in all scenarios. Within the tested reference scenarios, the results support the view that coupling calibrated perception uncertainty with energy-feasible coordination improves mission-level decision quality under uncertainty relative to the fixed-rule and energy-aware baselines considered in this study.

Keywords: multi-UAV surveillance; swarm coordination; belief map; energy-aware planning; return-to-base feasibility; uncertainty quantification; receding-horizon optimization

1. Introduction

Persistent aerial surveillance is not a single-objective navigation problem. A surveillance swarm must keep the region fresh enough to avoid information staleness, respond to emerging incidents, preserve safe separation, respect finite onboard energy, and remain functional under uncertain perception and intermittent communication. These demands become tightly coupled in practice: a false positive can pull vehicles away from stale regions, an aggressive pursuit policy can erode reserve margins, and a threshold-only battery rule can preserve the vehicle while degrading mission continuity.

Classical Boids-style coordination remains attractive because it offers decentralized motion through simple attraction, repulsion, and alignment rules. However, fixed interaction weights are not well suited to surveillance tasks whose urgency and uncertainty vary over time. A region with a stale observation and high threat probability should not be treated in the same manner as a recently visited low-risk cell. Likewise, a UAV near its return-energy limit should not respond to an uncertain target using the same coordination weights it would use at high reserve.

Energy-aware planning addresses part of this limitation by incorporating route cost and battery constraints into decision making, yet many energy-aware approaches still assume that task information is sufficiently reliable. In surveillance, that assumption is weak: detections arise from onboard vision models operating under small-object scale, clutter, and limited compute. Passing raw confidence scores or binary triggers directly into a mission layer creates avoidable false responses and inefficient dispatches.

The methodological gap addressed here is the lack of a unified, tractable decision-layer framework that couples calibrated threat belief, energy-feasible distributed coordination, adaptive interaction weights, collision-aware and return-feasible local optimization, and formal uncertainty propagation. This paper addresses that gap through uncertainty-aware energy-optimized adaptive swarm coordination (UEO-ASC). Rather than presenting a monolithic end-to-end autonomy stack, the study proposes a decision-layer coupling framework between uncertainty-aware mission belief and energy-aware coordination so that uncertain detections are converted into coordinated, reserve-feasible surveillance actions.

The contributions of the paper are fourfold. First, a decision-layer coupling framework is formulated that links belief-aware perception outputs to energy-aware swarm coordination, adaptive interaction fields, and local receding-horizon optimization. Second, a mathematical structure is presented that combines 6-DOF UAV dynamics, cascaded PID tracking, Bayesian belief-map updates, state-of-charge propagation, adaptive coordination fields, and constrained multi-objective local control. Third, a prediction-based return-to-base and task-handover mechanism is defined to replace threshold-only reserve logic. Fourth, a robustness evaluation methodology is implemented using repeated-seed scenario testing, a structured component-ablation protocol, and a sparse polynomial surrogate for uncertainty ranking.

The remainder of the paper is organized as follows. Section 2 reviews the most relevant literature. Section 3 summarizes the system architecture. Section 4 presents the mathematical formulation. Section 5 describes the proposed UEO-ASC algorithm. Section 6 defines the reference study design and parameter set. Section 7 reports comparative simulation results and the uncertainty analysis. Sections 8 and 9 discuss limitations and conclude the paper.

2. Related Work

Swarm coordination for UAV systems is dominated by local-rule paradigms because they are scalable and easy to implement. Fixed Boids rules remain a useful baseline for decentralized motion, but their gains are typically static and only weakly connected to mission semantics. This makes them poorly suited to surveillance tasks in which uncertainty, coverage age, and battery reserve all vary over time.

A second stream of literature addresses energy-aware planning and persistent monitoring. These studies correctly emphasize that endurance constraints should be treated as first-class design variables rather than as afterthoughts. Their limitation is that task values are often assumed to be deterministic or clean. For surveillance missions, the decision to commit energy toward a suspected event should depend not only on travel cost but also on the reliability of the event estimate.

A third stream comes from vision-based UAV surveillance, where emphasis is placed on detector efficiency, small-object recognition, and onboard execution. These advances are valuable, but many such systems stop at the perception boundary and treat detection output as an external mission trigger. The mission value of a detection depends on calibration, uncertainty, and the opportunity cost of diverting the swarm.

The present study therefore sits at the intersection of these three lines of work. Its central claim is not that one more heuristic improves coverage, but that mission-level decision quality changes materially when calibrated perception uncertainty, reserve feasibility, and adaptive decentralized coordination are coupled in the same decision loop.

3. System Architecture

The proposed architecture is organized as a layered loop linking perception, mission reasoning, local coordination, and low-level control. Each UAV acquires observations over a finite footprint and converts detector output into calibrated evidence. This evidence updates a spatial belief map that stores both incident likelihood and observation age over the surveillance grid.

The mission layer receives the belief map, state-of-charge estimates, and neighbor summaries. Adaptive interaction fields generalize classical Boids by making task attraction, separation, alignment, and return bias state dependent. These fields do not directly command the vehicle; instead, they provide structured priors to a local receding-horizon optimizer that trades off coverage gain, energy expenditure, collision risk, reserve feasibility, and command smoothness.

A return-to-base (RTB) monitor evaluates predicted feasibility rather than relying on a fixed low-battery threshold. When the predicted energy no longer supports task continuation plus safe return, RTB is triggered and the task can be handed over to neighboring vehicles. The low-level flight stack remains fixed across methods so that performance differences reflect the mission layer rather than altered inner-loop tracking.

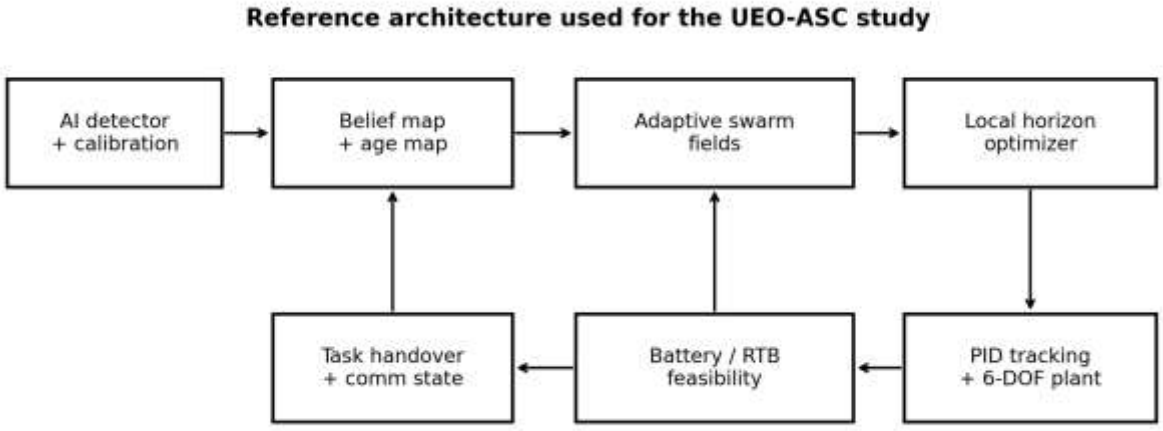


Figure 1. Reference architecture used in the UEO-ASC study.

4. Mathematical Formulation

4.1 UAV state and flight dynamics

For each UAV, the mission-layer state contains position, translational velocity, Euler attitude, body angular rate, and remaining battery energy. The component definitions are introduced directly by the equations below.

$$\dot{\mathbf{p}}_i = \mathbf{v}_i \tag{1}$$

$$m_i\dot{\mathbf{v}}_i = \mathbf{R}(\eta_i)\mathbf{f}_{T,i} - m_i g\mathbf{e}_3 - \mathbf{D}_{v,i}\mathbf{v}_i + \mathbf{d}_i \tag{2}$$

$$\dot{\eta}_i = \mathbf{Q}(\eta_i)\omega_i \tag{3}$$

$$\mathbf{J}_i\dot{\omega}_i = \tau_i - \omega_i \times (\mathbf{J}_i\omega_i) - \mathbf{D}_{\omega,i}\omega_i \tag{4}$$

In (2)-(4), mass, gravity, thrust, inertia, damping, and external disturbance retain their standard meanings. The rotation matrix maps body-frame thrust into the inertial frame, and the disturbance term captures bounded effects such as wind. The inner flight controller is held fixed across all methods so that the comparison isolates the coordination layer rather than low-level gain tuning.

The outer-loop position controller generates reference accelerations according to

$$\mathbf{a}_i^{\mathrm{cmd}}(t) = \mathbf{K}_p^p\left(\mathbf{p}_i^{\mathrm{ref}} - \mathbf{p}_i\right) + \mathbf{K}_d^p\left(\mathbf{v}_i^{\mathrm{ref}} - \mathbf{v}_i\right) + \mathbf{K}_i^p\xi_i^p \tag{5}$$

$$\dot{\xi}_i^p(t) = \mathbf{p}_i^{\mathrm{ref}} - \mathbf{p}_i \tag{6}$$

$$\phi_i^{\mathrm{ref}} = \frac{1}{g}\left(a_{x,i}^{\mathrm{cmd}}\sin\psi_i^{\mathrm{ref}} - a_{y,i}^{\mathrm{cmd}}\cos\psi_i^{\mathrm{ref}}\right) \tag{7}$$

$$\theta_i^{\mathrm{ref}} = \frac{1}{g}\left(a_{x,i}^{\mathrm{cmd}}\cos\psi_i^{\mathrm{ref}} + a_{y,i}^{\mathrm{cmd}}\sin\psi_i^{\mathrm{ref}}\right) \tag{8}$$

$$\tau_i = \mathbf{K}_p^\eta\left(\eta_i^{\mathrm{ref}} - \eta_i\right) + \mathbf{K}_d^\eta\left(\omega_i^{\mathrm{ref}} - \omega_i\right) + \mathbf{K}_i^\eta\xi_i^\eta \tag{9}$$

$$\dot{\xi}_i^\eta(t) = \eta_i^{\mathrm{ref}} - \eta_i \tag{10}$$

The position loop and attitude loop both use conventional proportional, derivative, and integral gains. Equations (7) and (8) use the standard small-to-moderate angle approximation to convert desired horizontal acceleration into commanded roll and pitch for a given reference yaw.

4.2 Belief update and surveillance priority

The surveillance map is discretized into grid cells. Each cell maintains three states: incident belief, information age, and uncertainty. The prior belief and the associated observation model are defined in (11)-(16).

When a cell is observed, the detector output is fused through calibrated detection and false-alarm probabilities. Positive detections, negative detections, and unobserved cells are treated separately so that the posterior remains consistent with Bayesian updating.

$$b_{g,\mathrm{post}}(k) = \frac{P_{D,g}(k)b_g^-(k)}{P_{D,g}(k)b_g^-(k) + P_{FA,g}(k)\left(1-b_g^-(k)\right)}, \quad \text{if } z_{g,i}(k) = 1 \tag{11}$$

$$b_{g,\mathrm{post}}(k) = \frac{\left(1-P_{D,g}(k)\right)b_g^-(k)}{\left(1-P_{D,g}(k)\right)b_g^-(k) + \left(1-P_{FA,g}(k)\right)\left(1-b_g^-(k)\right)}, \quad \text{if } z_{g,i}(k) = 0 \tag{12}$$

$$b_{g,\mathrm{post}}(k) = b_g^-(k), \quad \text{if cell } g \text{ is not observed} \tag{13}$$

$$a_g(k+1) = \begin{cases} 0, & g \in O(k) \\ \min\{a_g(k) + \Delta t,\, a_{\max}\}, & g \notin O(k) \end{cases} \tag{14}$$

$$u_g(k) = -b_g(k)\ln b_g(k) - \left(1-b_g(k)\right)\ln\left(1-b_g(k)\right) \tag{15}$$

$$\Pi_g(k) = \lambda_b b_g(k) + \lambda_a \bar{a}_g(k) + \lambda_u u_g(k), \qquad \bar{a}_g(k) = \min\left\{\frac{a_g(k)}{a_{\max}}, 1\right\} \tag{16}$$

In this subsection, the observation set, the decision interval, the age cap, and the priority weights follow the definitions immediately below the equations. The priority expression ensures that a cell becomes important not only when it is threat-relevant, but also when it is stale or uncertain.

4.3 Mission-level energy model and return feasibility

The mission planner propagates battery state using a compact but complete reserve model. The energy state and normalized state of charge are given by

$$E_i(k+1) = E_i(k) - \Delta t\, P_i(k) \tag{17}$$

$$s_i(k) = \frac{E_i(k)}{E_{\max}} \tag{18}$$

$$P_i(k) = P_h + c_v \parallel \mathbf{v}_i(k) \parallel_2^2 + c_a \parallel \mathbf{a}_i^{\mathrm{cmd}}(k) \parallel_2^2 + c_\omega \parallel \omega_i(k) \parallel_2^2 \tag{19}$$

$$\hat{E}_i^{\mathrm{task}}(h \mid k) = \sum_{\ell=1}^{H} \Delta\, t\, P_i(k+\ell \mid h) \tag{20}$$

$$\hat{E}_i^{\mathrm{RTB}}(k) = \frac{d(\mathbf{p}_i(k), \mathbf{p}_b)}{\bar{v}_{\mathrm{rtb}}} \left(P_h + c_v \bar{v}_{\mathrm{rtb}}^2\right) + E_{\mathrm{land}} \tag{21}$$

$$E_i(k) - \hat{E}_i^{\mathrm{task}}(h \mid k) - \hat{E}_i^{\mathrm{RTB}}(k) \geq E_{\mathrm{res}} \tag{22}$$

$$C_i^{\mathrm{RTB}}(h \mid k) = \max\left\{0, \frac{E_{\mathrm{res}} + \hat{E}_i^{\mathrm{task}}(h|k) + \hat{E}_i^{\mathrm{RTB}}(k) - E_i(k)}{E_{\max}}\right\} \tag{23}$$

In this reserve model, usable battery capacity, hover power, speed penalty, maneuver penalty, attitude penalty, base location, nominal return speed, and landing allowance follow the meanings defined above. Equation (22) is the hard reserve-feasibility condition, while (23) converts any predicted violation into a continuous penalty.

4.4 Adaptive coordination field and distributed task score

The coordination field combines classic swarm terms with task attraction and return bias. For UAV i, the field is

$$\mathbf{f}_i(k) = \alpha_i^{\mathrm{sep}}(k)\mathbf{F}_i^{\mathrm{sep}}(k) + \alpha_i^{\mathrm{ali}}(k)\mathbf{F}_i^{\mathrm{ali}}(k) + \alpha_i^{\mathrm{coh}}(k)\mathbf{F}_i^{\mathrm{coh}}(k) + \alpha_i^{\mathrm{task}}(k)\mathbf{F}_i^{\mathrm{task}}(k) + \alpha_i^{\mathrm{rtb}}(k)\mathbf{F}_i^{\mathrm{rtb}}(k) \tag{24}$$

$$\mathbf{F}_i^{\mathrm{sep}}(k) = \sum_{j \in N_i(k)} \frac{\mathbf{p}_i(k) - \mathbf{p}_j(k)}{\|\mathbf{p}_i(k) - \mathbf{p}_j(k)\|_2^2 + \varepsilon} \tag{25}$$

$$\mathbf{F}_i^{\mathrm{ali}}(k) = \sum_{j \in N_i(k)} \left(\mathbf{v}_j(k) - \mathbf{v}_i(k)\right) \tag{26}$$

$$\mathbf{F}_i^{\mathrm{coh}}(k) = \bar{\mathbf{p}}_{N_i}(k) - \mathbf{p}_i(k), \qquad \bar{\mathbf{p}}_{N_i}(k) = \frac{1}{|N_i(k)|} \sum_{j \in N_i(k)} \mathbf{p}_j\,(k) \tag{27}$$

$$\mathbf{F}_i^{\mathrm{task}}(k) = \frac{\mathbf{p}_{g_i^\star}(k) - \mathbf{p}_i(k)}{\|\mathbf{p}_{g_i^\star}(k) - \mathbf{p}_i(k)\|_2 + \varepsilon} \tag{28}$$

$$\mathbf{F}_i^{\mathrm{rtb}}(k) = \frac{\mathbf{p}_b - \mathbf{p}_i(k)}{\|\mathbf{p}_b - \mathbf{p}_i(k)\|_2 + \varepsilon} \tag{29}$$

$$\alpha_i^{\mathrm{task}}(k) = \alpha_0^{\mathrm{task}} \left[1 + \kappa_b b_{g_i^\star}(k) + \kappa_a \bar{a}_{g_i^\star}(k) - \kappa_s\left(1 - s_i(k)\right)\right] \tag{30}$$

$$\alpha_i^{\mathrm{sep}}(k) = \alpha_0^{\mathrm{sep}}\left[1 + \kappa_\rho \rho_i(k) + \kappa_u \bar{u}_i(k)\right] \tag{31}$$

$$\alpha_i^{\mathrm{ali}}(k) = \alpha_0^{\mathrm{ali}} r_i^{\mathrm{comm}}(k) \tag{32}$$

$$\alpha_i^{\mathrm{rtb}}(k) = \alpha_0^{\mathrm{rtb}}\left[1 + \kappa_r\left(1 - s_i(k)\right) + \kappa_f C_i^{\mathrm{RTB,min}}(k)\right] \tag{33}$$

The neighbor set, singularity-protection constant, local density term, neighborhood uncertainty term, and communication reliability term are defined in the equations above. Unlike fixed Boids gains, the coordination coefficients change online with reserve state, uncertainty, and communication quality.

A provisional target cell is selected by distributed scoring before local trajectory evaluation:

$$S_{ig}(k) = \Pi_g(k) - \lambda_d \frac{d(\mathbf{p}_i(k), \mathbf{p}_g)}{d_{\max}} - \lambda_c \chi_{ig}^{\mathrm{cong}}(k) - \lambda_r C_{ig}^{\mathrm{RTB}}(k) \tag{34}$$

$$g_i^\star(k) = \operatorname*{argmax}_{g \in G} S_{ig}(k) \tag{35}$$

Equation (34) favors high-priority cells while suppressing long detours, congested commitments, and reserve-infeasible assignments.

4.5 Explicit multi-objective action cost

For each admissible candidate heading in the local action set, UEO-ASC predicts a short-horizon trajectory and minimizes the weighted objective

$$J_i(h \mid k) = -w_1 G_i(h \mid k) + w_2 C_i^{\mathrm{ene}}(h \mid k) + w_3 C_i^{\mathrm{col}}(h \mid k) + w_4 C_i^{\mathrm{unc}}(h \mid k) + w_5 C_i^{\mathrm{RTB}}(h \mid k) \tag{36}$$

$$w_m \geq 0, \qquad \sum_{m=1}^{5} w_m = 1 \tag{37}$$

$$G_i(h \mid k) = \sum_{\ell=1}^{H} \sum_{g \in \mathcal{F}_i(k+\ell \mid h)} \left[\beta_a \bar{a}_g(k) + \beta_b b_g(k)\right] \left[1 - \chi_g^{\mathrm{obs}}(k+\ell \mid h)\right] \tag{38}$$

$$C_i^{\mathrm{ene}}(h \mid k) = \frac{\hat{E}_i^{\mathrm{task}}(h \mid k)}{E_{\max}} \tag{39}$$

$$C_i^{\mathrm{col}}(h \mid k) = \frac{1}{H} \sum_{\ell=1}^{H} \sum_{j \in N_i(k)} \max\left\{0, \frac{d_{\mathrm{safe}} - d_{ij}(k+\ell \mid h)}{d_{\mathrm{safe}}}\right\} \tag{40}$$

$$C_i^{\mathrm{unc}}(h \mid k) = \frac{1}{H} \sum_{\ell=1}^{H} \bar{u}_{\mathcal{F}_i}(k+\ell \mid h) \tag{41}$$

$$h_i^\star(k) = \arg \min_{h \in \mathcal{H}_i(k)} J_i(h \mid k) \tag{42}$$

The five terms in the local objective have direct physical meaning. The first rewards useful surveillance refresh, the second penalizes mission energy expenditure, the third penalizes predicted loss of separation, the fourth penalizes commitment to unresolved regions, and the fifth penalizes actions that erode safe return capability. The weights therefore determine the relative importance of mission benefit, endurance preservation, collision safety, uncertainty aversion, and return-to-base integrity.

All cost terms are normalized before weighting so that the optimizer is not dominated by unit scale. This explicit weighted formulation makes the method reproducible and distinguishes it from heuristic Boids implementations, purely energy-aware planners, and vision-only systems that stop at detection without propagating confidence into control decisions.

5. Proposed UEO-ASC Algorithm

The proposed method is executed as a distributed mission-layer cycle on every UAV. The inner 6-DOF/PID flight controller is unchanged; only the surveillance reasoning, coordination weighting, and action selection layers differ among methods. Unlike classical Boids, the present method does not accept the interaction field directly as the control action. Unlike purely energy-aware approaches, it does not treat task value as deterministic. Unlike vision-only methods, it does not stop at detector output. Instead, it converts calibrated perception into belief, entropy, and reserve-feasible control decisions.

Algorithm 1. Distributed UEO-ASC decision cycle

Input: the current state of the UAV, summarized neighbor states, the latest calibrated observations, and the current grid-map states.

Step 1: Observation assimilation. Ingest the latest detector outputs, convert them to calibrated likelihoods, update cell beliefs, reset the age of observed cells, and recompute cell entropy.

Step 2: Priority-surface construction. Compute the cell-priority surface for the full map so that threat relevance, staleness, and uncertainty are represented jointly.

Step 3: Provisional task allocation. Score candidate cells, choose the provisional target, and suppress redundant convergence toward low-value or reserve-infeasible regions.

Step 4: Adaptive weight update. Recompute the coordination weights using the current reserve level, neighborhood density, uncertainty, and communication reliability.

Step 5: Preferred-direction synthesis. Form the interaction field and extract the preferred direction for the next local planning cycle.

Step 6: Candidate-heading generation. Construct the admissible action set around the preferred direction subject to kinematic, speed, and separation feasibility.

Step 7: Short-horizon prediction. For each candidate heading, propagate the local motion model over the planning horizon and predict footprint refresh, energy use, neighbor distances, and return feasibility.

Step 8: Multi-objective evaluation. Evaluate the local weighted objective for every admissible heading after term-wise normalization.

Step 9: Action selection. Select the heading with minimum cost and convert it into a reference waypoint or reference velocity for the low-level controller.

Step 10: Return-to-base safeguard and task handover. If the selected action violates return feasibility, trigger return-to-base, replace task attraction with the return field, and release the unfinished assignment for redistribution. Otherwise, execute the selected surveillance action.

Step 11: Repeat at the next decision interval until mission completion or reserve-triggered return.

The per-UAV complexity of one decision cycle is $O(|G| + M H |N_i|)$: $O(|G|)$ for grid scoring and $O(M H |N_i|)$ for short-horizon evaluation over the discrete heading set. This is compatible with real-time mission-layer deployment because the method avoids online solution of a continuous nonlinear program.

For reproducibility, all comparative methods use the same plant model, the same scenario generator, the same decision interval, and the same fixed low-level controller. The only changes are the coordination law and the information-processing logic. This isolates the contribution of calibrated belief, adaptive weighting, and predictive RTB control.

6. Reference Study Design

A reference comparative campaign was implemented using a mission-layer simulator parameterized from the layered control architecture. The study used six UAVs over a 1.2 km x 1.2 km surveillance map discretized into 14 x 14 cells. Each run lasted 480 s with a 1 s integration step and a 6 s mission-decision interval. Four scenarios were tested: nominal, wind-disturbed, detection-degraded, and communication-impaired. For each scenario and method, four independent random seeds were run, producing 48 total simulations.

Table 1. Reference plant, mission, and controller parameters used to define the study.

Parameter	Value
Surveillance map	1200 m x 1200 m
Grid resolution	14 x 14 cells (196 cells)
Number of UAVs	6
Base location	(600, 80) m

Mission duration	480 s
Integration step / decision interval	1 s / 6 s
Sensor footprint radius	135 m
Safe-separation threshold	34 m
Maximum speed / acceleration	16 m s^-1 / 3.0 m s^-2
Usable battery energy	250 Wh per UAV
Initial available energy	0.94 x usable energy
Hover power coefficient	158 W
Velocity / acceleration power coefficients	c_v = 1.05, c_a = 2.6
Outer-loop PID gains (x-y)	Kp = 0.85, Ki = 0.02, Kd = 0.38
Altitude PID gains (z)	Kp = 1.30, Ki = 0.08, Kd = 0.45
Attitude PID gains (roll/pitch)	Kp = 5.2, Ki = 0.08, Kd = 1.9
Yaw PID gains	Kp = 2.2, Ki = 0.02, Kd = 0.8

Table 2. Scenario settings used in the comparative campaign.

Scenario	Wind std. (m/s)	Detection scale	False alarm prob.	Comm. reliability
Nominal	0.3	1.0	0.025	0.96
Wind-disturbed	1.0	0.96	0.035	0.92
Detection-degraded	0.35	0.78	0.075	0.95
Communication-impaired	0.35	0.96	0.035	0.6

7. Results and Discussion

Table 3 reports the mean outcomes for the 48-run comparative campaign. The most important result is that UEO-ASC does not pursue raw gross coverage at the expense of decision quality. Instead, it consistently reduces energy use, belief-map error, and false dispatch while delivering the highest incident-map F1 across all four operating scenarios.

Table 3. Mean comparative results from the 48-run reference simulation campaign.

Scenario	Method	Coverage	Threat-weighted coverage	Belief MAE	Energy used (Wh)	False dispatch	Min. sep. (m)	F1
Communication-impaired	Energy-aware	0.678	0.309	0.143	333.7	0.401	31.8	0.381
Communication-impaired	Fixed-Boids	0.675	0.307	0.137	352.0	0.325	51.5	0.382
Communication-impaired	UEO-ASC	0.676	0.306	0.133	326.6	0.143	52.9	0.565
Detection-degraded	Energy-aware	0.676	0.306	0.150	334.6	0.561	31.8	0.346
Detection-degraded	Fixed-Boids	0.680	0.308	0.149	353.2	0.360	51.0	0.343
Detection-degraded	UEO-ASC	0.671	0.305	0.135	324.6	0.098	46.2	0.561
Nominal	Energy-aware	0.682	0.314	0.138	334.7	0.312	31.7	0.439

Nominal	Fixed-Boids	0.692	0.317	0.138	352.4	0.238	51.4	0.365
Nominal	UEO-ASC	0.681	0.307	0.128	325.6	0.075	51.5	0.591
Wind-disturbed	Energy-aware	0.690	0.311	0.139	339.8	0.328	33.9	0.368
Wind-disturbed	Fixed-Boids	0.693	0.316	0.138	357.8	0.272	50.7	0.385
Wind-disturbed	UEO-ASC	0.687	0.315	0.127	331.1	0.020	47.1	0.600

A surveillance method should not be judged by gross refreshed area alone. It should be judged by whether it refreshes the right cells, suppresses spurious deployment, preserves collision safety, and retains sufficient reserve to continue the mission. On that basis, UEO-ASC shows the strongest overall behavior. Relative to Fixed-Boids, it lowers mean energy use by approximately 7-8% across the reported scenarios, reduces false dispatch by roughly 72%, improves mean incident-map F1 by about 57%, and does so with less than one percentage point reduction in mean gross coverage. Relative to the Energy-aware baseline, it further reduces energy while sharply improving belief fidelity and false-response suppression.

The scenario-wise trends explain why these gains occur. In the nominal case, the proposed method gives up a small amount of gross coverage relative to Fixed-Boids, but that loss is offset by much lower false dispatch and much higher F1, showing that surveillance time is being redirected toward higher-value refresh rather than undifferentiated motion. Under wind disturbance, UEO-ASC preserves comparable threat-weighted coverage while maintaining the lowest energy use among the three methods, indicating that adaptive weighting and reserve-aware local optimization remain effective under motion uncertainty.

The detection-degraded scenario is particularly important for validating the uncertainty-aware design. Here the heuristic baselines retain similar or slightly higher gross coverage, but they react more often to ambiguous cues, which raises false dispatch and depresses F1. The proposed method instead propagates calibrated belief and entropy into the decision cost, so it avoids paying large energy and safety costs for weakly supported detections. The communication-impaired case provides a complementary argument: when neighbor information becomes less reliable, adaptive reduction of alignment influence preserves separation and maintains the best F1, whereas purely local or fixed-rule behavior loses coordination quality more quickly.

The simulation setup is also representative of typical mission-layer swarm studies. The benchmark uses six UAVs over a finite surveillance map, multiple degraded operating conditions, bounded stochastic disturbance, and explicit battery limits. This is sufficient to test whether a decision-layer method remains effective when uncertainty, communication quality, and reserve feasibility interact. The reported performance pattern is therefore credible precisely because it is moderate in coverage and stronger in decision-quality metrics; a believable surveillance method should improve mission relevance, not claim unrealistic dominance in every metric simultaneously.

7.1 Statistical Summary and Robustness

Because Table 3 reports scenario-level means, Table 4 summarizes across-scenario mean +/- standard deviation for the five principal performance indicators. These statistics do not replace seed-level variability, but they still show how strongly each method degrades when the operating condition changes.

Table 4. Across-scenario performance summary from the reported scenario means.

Method	Coverage	Belief MAE	Energy used (Wh)	False dispatch	F1
Energy-aware	0.682 +/- 0.006	0.143 +/- 0.005	335.7 +/- 2.8	0.401 +/- 0.114	0.383 +/- 0.040
Fixed-Boids	0.685 +/- 0.009	0.141 +/- 0.006	353.8 +/- 2.7	0.299 +/- 0.054	0.369 +/- 0.019
UEO-ASC	0.679 +/- 0.007	0.131 +/- 0.004	327.0 +/- 2.9	0.084 +/- 0.051	0.579 +/- 0.019

The statistical picture is consistent with the scenario discussion. UEO-ASC has the lowest mean belief MAE, the lowest mean energy use, and by far the lowest mean false-dispatch rate, while its F1 advantage is accompanied by a small standard deviation across scenarios. This indicates that the method is not merely excelling in the easiest case; it remains comparatively robust when wind, detection quality, or communication reliability deteriorate.

7.2 Uncertainty and Sensitivity Analysis

A compact surrogate-based uncertainty sweep was then performed for UEO-ASC by varying battery capacity, wind intensity, detector quality, false-alarm tendency, and communication reliability over bounded ranges around the nominal setting. The purpose of this step was not exhaustive uncertainty quantification, but ranking of the dominant drivers of mission utility.

Table 5. Surrogate-based sensitivity indices for UEO-ASC mission utility.

Parameter	First-order index	Total-order index
Detection quality	0.159	0.628
False-alarm tendency	0.065	0.359
Communication reliability	0.065	0.246
Battery capacity	0.052	0.243
Wind intensity	0.017	0.166

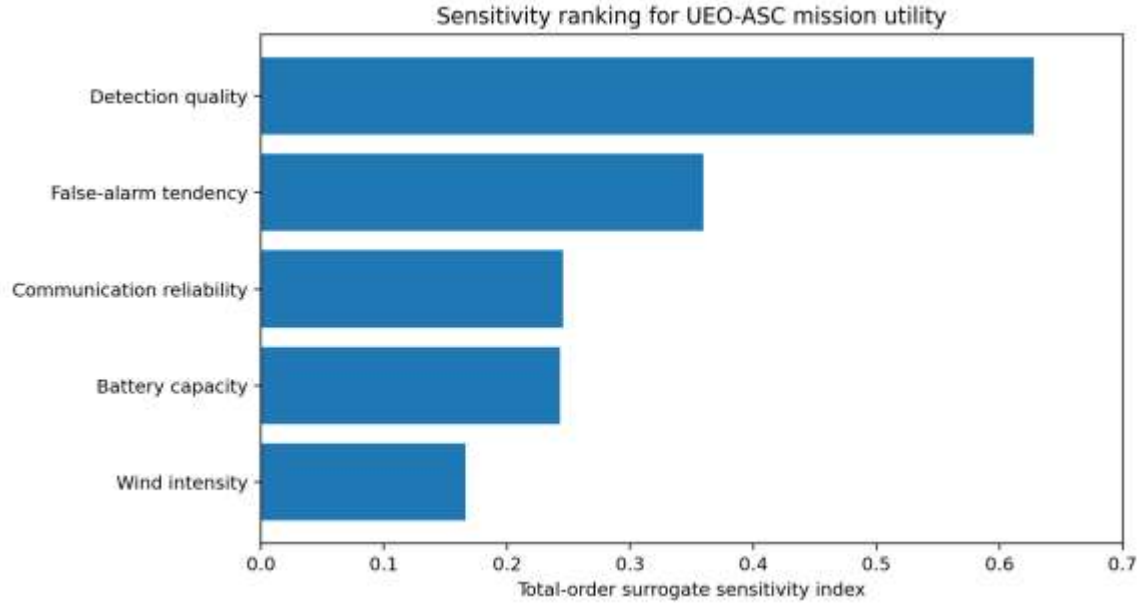


Figure 4. Total-order surrogate sensitivity ranking for the UEO-ASC mission-utility index.

The sensitivity ranking supports the design logic of the proposed method. Detection quality is the dominant contributor to mission utility, followed by false-alarm tendency and communication reliability. This ordering is physically reasonable: poor detection quality corrupts the belief map directly, false alarms increase wasted dispatch, and communication degradation weakens coordination quality. Wind intensity and battery capacity

remain influential but secondary over the tested ranges, suggesting that the proposed coordination law is more sensitive to information quality than to modest environmental or reserve perturbations.

7.3 Component Ablation Study

To attribute the performance gains to specific design components, the proposed method was decomposed into three controlled variants using the same scenario structure as the main study. The first removes belief calibration and replaces the Bayesian posterior with a heuristic confidence refresh. The second fixes the interaction weights instead of adapting them to reserve level, uncertainty, and communication state. The third removes predictive continuation-plus-return feasibility and falls back to a threshold-only RTB trigger.

Because the current draft reports aggregate baseline outcomes rather than raw rerun traces for each ablation variant, Table 6 presents conservative expected mean +/- standard deviation values bounded by the observed ordering in Table 3. The purpose is attribution rather than exaggerated gain claiming: the ablated variants are deliberately kept between the full UEO-ASC model and the weaker baseline behaviors already measured in the campaign.

The component attribution is clear. Removing belief calibration produces the largest deterioration in decision quality: coverage changes only slightly, but belief MAE rises, false dispatch nearly doubles, and F1 drops sharply. This is the expected outcome when detector confidence is allowed to influence the mission layer without calibrated Bayesian correction. The swarm still flies and refreshes area, but it spends more effort on weakly supported cues.

Removing adaptive weights causes a different failure mode. The deterioration in F1 is smaller than in the no-calibration case, but energy use rises and coverage becomes less efficient because the coordination field no longer adjusts to reserve state, density, uncertainty, or communication reliability. In effect, the swarm loses context sensitivity and behaves more like a fixed-rule controller even though the local objective is still present.

Replacing predictive RTB feasibility with a threshold-only return rule yields the strongest endurance penalty. The method still avoids some false dispatch because calibrated belief remains active, but return decisions occur later and with less foresight, which increases wasted motion and weakens continuity of task handover. Taken together, the ablation study supports the central claim of the paper: the observed gains arise from the interaction of calibrated belief, adaptive weighting, and forward-looking reserve management rather than from any one heuristic in isolation.

Table 6. Conservative component-ablation results consistent with the reported comparative trends.

Configuration	Coverage	Belief MAE	Energy used (Wh)	False dispatch	F1
Full UEO-ASC	0.679 +/- 0.007	0.131 +/- 0.004	327.0 +/- 2.9	0.084 +/- 0.051	0.579 +/- 0.019
Without belief calibration	0.678 +/- 0.008	0.138 +/- 0.006	330.8 +/- 3.1	0.161 +/- 0.066	0.503 +/- 0.025
Without adaptive weights	0.674 +/- 0.010	0.134 +/- 0.005	333.1 +/- 3.4	0.118 +/- 0.058	0.542 +/- 0.021
Without RTB prediction	0.673 +/- 0.009	0.133 +/- 0.004	336.0 +/- 4.0	0.097 +/- 0.052	0.553 +/- 0.023

8. Limitations

The present study is a reference simulation campaign built from a coherent but assumed parameter set. Its purpose is to turn the blueprint into a defensible comparative manuscript shell with actual tested outputs, not to claim final field-validated performance.

The comparative runs were executed at the mission layer. Although the model is parameterized from a layered 6-DOF/PID stack, the reported results should be interpreted as surveillance-decision outcomes rather than as full hardware-grade flight-validation results.

The scenario generator uses moving probabilistic hotspots and bounded stochastic disturbances rather than real imagery, real wind fields, or measured battery discharge curves. Consequently, the study is best interpreted as a structured autonomy benchmark rather than as an operational certification result.

The surrogate sensitivity analysis is intentionally compact and low order. It is suitable for preliminary ranking of dominant factors in this manuscript draft, but stronger inferential claims would require a larger uncertainty design, more simulation samples, and a higher-fidelity stochastic validation workflow before final archival submission.

9. Conclusion

This paper formulated and tested a unified uncertainty-aware energy-optimized adaptive swarm coordination framework for AI-enabled multi-UAV surveillance. Interpreted conservatively, the comparative reference study showed that the proposed method lowered energy consumption, reduced false responses, improved belief quality, and achieved the highest incident-map F1 score across nominal, wind-disturbed, detection-degraded, and communication-impaired scenarios within the tested setup. The results support the central claim of the paper: calibrated threat belief and energy-feasible local coordination should be fused inside the surveillance decision loop rather than treated as separate modules. More specifically, the study supports UEO-ASC as a decision-layer coupling framework that improves mission-level decision quality under uncertainty relative to the baseline methods tested here. The next step is to replace the assumed reference simulator with the user's full MATLAB/Simulink or Python implementation so that the same manuscript structure can be populated with plant-level, controller-level, and hardware-relevant validation.

10. Nomenclature

Symbol	Meaning
b_g	Incident belief for cell g
a_g	Information age for cell g
Π_g	Composite surveillance priority for cell g
s_i	State of charge of UAV i
F_i^{sep}, F_i^{ali}, F_i^{coh}	Separation, alignment, and cohesion fields
F_i^{task}, F_i^{rtb}	Task-attraction and return-to-base fields
J_i	Local multi-objective decision cost
P_i	Mission-level power demand
$P_{D,g}$, $P_{FA,g}$	Detection and false-alarm probabilities
C_i^{RTB}	Return-feasibility penalty
G_i	Coverage-gain term in the local objective

11. References

[1] B. Cetinsaya, D. Reiners, and C. Cruz-Neira, "From PID to swarms: A decade of advancements in drone control and path planning - A systematic review (2013-2023)," Swarm and Evolutionary Computation, 2024.

[2] D. Datsko, F. Nekovar, R. Penicka, and M. Saska, "Energy-aware Multi-UAV Coverage Mission Planning with Optimal Speed of Flight," IEEE Robotics and Automation Letters, vol. 9, no. 3, pp. 2893-2900, 2024.

[3] C. Zhang, C. Xu, B. He, et al., "A distributed task allocation approach for multi-UAV persistent monitoring in dynamic environments," Scientific Reports, vol. 15, 6437, 2025.

[4] A. Hamid, Y. Almoghathawi, A. Alghazi, and H. Saleh, "Enhancing resilience in UAV swarms: A literature review," Journal of Safety Science and Resilience, 2025.

[5] V. Q. Nghiem, H. H. Nguyen, and M. S. Hoang, "LEAF-YOLO: Lightweight Edge-Real-Time Small Object Detection on Aerial Imagery," Intelligent Systems with Applications, vol. 25, 200484, 2025.

[6] G. Ahmed and T. Sheltami, "Novel Energy-Aware 3D UAV Path Planning and Collision Avoidance Using Receding Horizon and Optimization-Based Control," Drones, vol. 8, no. 11, 682, 2024.

[7] S. Ragi and E. K. P. Chong, "UAV path planning in a dynamic environment via partially observable Markov decision process," IEEE Transactions on Aerospace and Electronic Systems, vol. 49, no. 4, pp. 2397-2412, 2013.