Cognitive Radar for Predictive UAV Surveillance: A Review of Sensing, Tracking, Learning, Re-Acquisition, And Threat Assessment

Ms. Suvarna Sahebrao Patil[1], Dr. Shivleela Mudda[2]
[1]PG Student, Department of Electronics and Telecommunication Engineering, M. S. Bidve Engineering College, Latur, Maharashtra, India
[2]Associate Professor, Department of Electronics and Telecommunication Engineering, M. S. Bidve Engineering College, Latur, Maharashtra, India

***Abstract*—Small unmanned aerial vehicles (UAVs) are difficult radar targets because low radar cross section, low-altitude clutter, multipath, rapid maneuvering, and intermittent visibility jointly degrade detection and track continuity. Cognitive radar is attractive in this setting because it closes the perception-action loop: the radar does not only estimate a target state, but can adapt sensing resources according to uncertainty, predicted motion, and mission risk. This paper presents a PRISMA-informed systematic review and critical synthesis of cognitive radar methods for predictive UAV surveillance through August 2026. The evidence is organized across six coupled functions: FMCW/mmWave sensing and micro-Doppler perception; uncertainty-aware single- and multi-target tracking; learning-based classification and trajectory prediction; adaptive radar resource management; autonomous target re-acquisition; and interpretable threat assessment. Recent 4-D radar studies demonstrate the value of joint detection-tracking and transformer-based trajectory modeling, while contemporary deep-reinforcement-learning work demonstrates increasingly capable time and task allocation. However, the literature remains fragmented: high classification accuracy is often reported on private or sensor-specific datasets, learned prediction is not consistently calibrated against tracking uncertainty, recovery after target loss is rarely treated as a primary metric, and threat scores lack a shared ontology. The review therefore proposes an integrated reference architecture and a benchmark framework spanning detection, tracking, recovery, risk, and embedded latency. The main conclusion is that publishable progress now depends less on adding isolated deep models and more on reproducible end-to-end evaluation under realistic clutter, occlusion, multi-target ambiguity, computational constraints, and explicit uncertainty.**

***Index Terms*—cognitive radar, UAV surveillance, FMCW radar, micro-Doppler, multi-target tracking, trajectory prediction, reinforcement learning, target re-acquisition, threat assessment, 4-D radar.**

I. INTRODUCTION

The rapid diffusion of low-cost multirotor and fixed-wing UAVs has changed the surveillance problem from one dominated by conventional aircraft to one in which small, slow, low-altitude, and highly maneuverable targets may appear close to terrain, buildings, vegetation, and other strong reflectors. Radar remains important because it directly provides range and radial-motion information and can operate independently of daylight, but the useful radar return from a small UAV can be intermittent and highly aspect dependent. Rotor blades introduce micro-Doppler components that are valuable for recognition, yet the same signatures vary with rotor speed, payload, observation angle, waveform, dwell time, and propagation environment [13]-[21].

A conventional surveillance chain detects a measurement, initiates a track, estimates kinematics, classifies the object, and then forwards the result to an operator. The weakness of this feed-forward view is that each stage consumes information without necessarily influencing the next sensing action. When a track becomes uncertain, the radar may continue its fixed scan even though a shorter revisit or a different beam allocation could materially improve recovery. Cognitive radar was introduced as a feedback-oriented paradigm in which the sensing system learns from the environment, maintains memory, and adapts the transmitted or scheduled action to the current objective [3], [4]. The key implication for UAV surveillance is that uncertainty, predicted motion, and threat can become inputs to radar control rather than being passive outputs of signal processing. The literature has advanced rapidly but unevenly. Classical Bayesian tracking provides transparent uncertainty propagation and mature tools for data association [29]-[34]. Deep models now operate on range-Doppler maps, micro-Doppler spectrograms, target tracks, and 4-D radar streams [16]-[28], [35], [36]. Reinforcement learning (RL) and constrained optimization have become increasingly credible for radar resource management [6]-[12]. At the decision layer, probabilistic and fuzzy methods are being used to represent UAV threat in multi-asset environments [37]-[40]. What remains relatively rare is an end-to-end study that connects all of these pieces and evaluates the coupling among them. This review therefore asks a systems question rather than a single-algorithm question: what combination of sensing, estimation, prediction, cognitive control, recovery, and risk evaluation is supported by current evidence, and what must be demonstrated before a cognitive radar architecture can be considered deployable for predictive UAV surveillance? The paper extends the supplied review framework by refreshing the literature through August 2026, tightening the distinction between evidence and proposal, adding recent 4-D radar and constrained-RL studies, and converting recurring weaknesses into a concrete experimental checklist.

II. REVIEW SCOPE, QUESTIONS, AND METHOD

A. Research Questions

Six research questions structure the synthesis. RQ1 examines which radar representations are most useful for small-UAV detection and discrimination. RQ2 examines how tracking methods preserve identity and uncertainty in clutter and multi-target scenes. RQ3 asks how learned models support trajectory prediction without discarding uncertainty. RQ4 examines how cognitive radar allocates waveform, dwell, revisit, beam, and task resources. RQ5 asks how a radar should recover a temporarily lost target. RQ6 examines how radar-derived behavior should be converted into an interpretable threat assessment rather than an opaque alert label.

TABLE I. REVIEW QUESTIONS AND REPRESENTATIVE SEARCH CONCEPTS

RQ	Focus	Representative concepts
RQ1	Sensing / perception	FMCW, mm Wave, 4-D MIMO, range-Doppler, micro-Doppler
RQ2	Tracking	EKF/UKF, IMM, PDA/JPDA, MHT, PHD/GLMB, OSPA
RQ3	Prediction	LSTM/GRU, transformer, track features, hybrid estimation-learning
RQ4	Cognitive control	resource management, dwell/revisit, beam, RL, POMDP
RQ5	Recovery	missed detections, occlusion, re-acquisition, gating, track-before-detect
RQ6	Threat	risk, intent, approach geometry, multi-asset scoring, explainability

B. Search and Screening Strategy

The review uses PRISMA-informed reporting principles [1] and the reproducibility logic of systematic engineering reviews [2]. A structured search was refreshed on 12 August 2026 using publisher and index interfaces for IEEE Xplore, IET/Wiley, Elsevier/ScienceDirect, SpringerLink, MDPI, and DOI/citation cross-checking. The principal time window is 2006-2026 because the modern cognitive-radar formulation dates from 2006 [3]. Older work is retained where it defines still-current foundations such as probabilistic data association, multiple-hypothesis tracking, interacting multiple models, PHD/GLMB filtering, and OSPA metrics [29]-[34]. Candidate studies were included when radar was a central sensing modality and the paper supplied enough technical information to extract at least the input representation, algorithmic role, evaluation setting, and limitations. Optical-only, acoustic-only, and communication-only studies were excluded as primary evidence. Reviews were used to orient terminology, but performance claims in the synthesis were anchored to experimental or algorithmic primary studies whenever possible. Recent preprints were not treated as equivalent to peer-reviewed evidence when a journal or conference version was available. Because commercial indexes expose different query grammars and duplicate records, this manuscript does not invent a pooled search-hit count; instead, it reports the search logic and a traceable evidence set. This choice is more reproducible than a nominal PRISMA number that cannot be reconstructed from exported records.

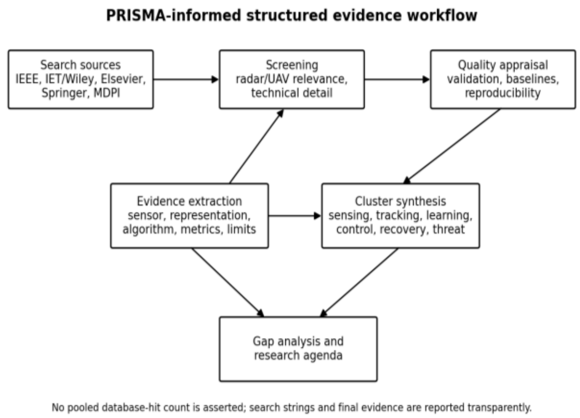

Fig. 1. PRISMA-informed structured workflow used to identify, appraise, cluster, and synthesize the evidence. The workflow emphasizes traceability without asserting unverifiable pooled database-hit counts.

C. Quality Appraisal

Each study was appraised against five dimensions: (1) clarity of target and environment; (2) transparency of radar representation and algorithm; (3) strength of validation, including real measurements and baselines; (4) reproducibility, including sensor and parameter reporting; and (5) limitation reporting, including false alarms, clutter, latency, memory, power, and generalization.

Evidence weight was reduced when evaluation relied only on simulation or a small private dataset without scenario variation. This distinction is important because a model can be internally accurate while still being poorly supported for outdoor surveillance.

III. RADAR SENSING AND SIGNAL REPRESENTATIONS

A. FMCW and 4-D Radar

FMCW radar is common in compact UAV surveillance because a linear frequency sweep supports simultaneous range and radial-velocity estimation.

For an ideal linear chirp with slope S and beat frequency f_b, a first-order range estimate is

$$R = c f_b / (2 S) \qquad (1)$$

where c is the speed of light. Radial velocity is obtained from Doppler frequency f_D as

$$v_r = \lambda f_D / 2 \qquad (2)$$

where λ is wavelength. Practical systems add angle estimation through antenna arrays and MIMO processing. Recent 4-D FMCW systems provide range, Doppler, azimuth, and elevation information, which is particularly valuable when a small target is embedded in clutter because spatial consistency can be exploited across frames rather than relying on a single thresholder detection [35], [36].

The representation made available by the hardware strongly constrains the achievable algorithms. Raw ADC samples enable customized FFT processing, clutter removal, CFAR, beamforming, and micro-Doppler analysis but increase bandwidth and computation. Range-Doppler and angle-Doppler maps preserve more structure than target lists while reducing the raw data volume. Commercial modules that expose only target lists can be excellent for kinematic prototypes, but they hide thresholding, confidence, and signal-quality information that would otherwise support track-before-detect or uncertainty-aware learning.

B. Micro-Doppler as a Discriminative Cue

Micro-Doppler arises when rotating or vibrating parts impose additional frequency modulation on the radar return [13]. For UAVs, rotor motion produces periodic or quasi-periodic structure that can help distinguish multirotor from birds and some background objects.

Early work demonstrated that such signatures can support small-UAV discrimination [14]-[17]. Subsequent CNN-based methods showed that time-frequency spectrograms and range-Doppler imagery can be learned directly [16], [19]. Nevertheless, the micro-Doppler signature is not an intrinsic, sensor-independent class label.
It changes with aspect, carrier frequency, coherent processing interval, rotor configuration, motion, SNR, and clutter.

A strong evaluation should therefore separate random train-test splits from cross-flight, cross-aspect, cross-location, and cross-UAV generalization.

TABLE II. REPRESENTATIVE RADAR-UAV PERCEPTION STUDIES

Study	Input / method	Main contribution	Key caution
Huizing et al. 2019 [16]	Micro-Doppler; deep classifier	Mini-UAV classification in cognitive-radar setting	Signature varies with geometry and data domain
Oh et al. 2019 [17]	FMCW micro-Doppler features	Systematic UAV signature analysis and classification	Controlled sensor/target conditions
Park et al. 2020 [18]	FMCW range-Doppler processing	Improves small moving-drone map concentration	Detection stage does not solve identity continuity
Rahman & Robertson 2020 [19]	Spectrogram CNN	Drone-bird discrimination from micro-Doppler	Dataset/sensor transfer remains a concern
Liu et al. 2021 [22]	Radar tracks; random forest	Uses motion behavior for bird/drone discrimination	Depends on track quality and scenario coverage
Yang et al. 2023 [23]	Radar echoes; transformer	Echoformer learns long-range echo relationships	Deep model demands representative training data
Janpangngern et al. 2025 [26]	High-resolution FMCW / SDR	Recent measurement-oriented small-UAV detection platform	Hardware-specific processing chain
Zhang & Song 2026 [27]	Dual-band micro-Doppler + Mamba	Modern sequence model with complementary bands	Performance is dataset-specific; dual-band cost is higher
Tong et al. 2026 [28]	Track features; TCN-Transformer + TimeGAN	Addresses scarce and imbalanced trajectory data	Synthetic augmentation must preserve real dynamics

IV. UNCERTAINTY-AWARE TRACKING AND MULTI-TARGET CONTINUITY

A. Recursive State Estimation

Tracking remains the mathematical backbone of predictive surveillance because it represents both the estimated state and its uncertainty. For a linear Gaussian model, the recursive prediction can be written

$$\hat{x}_k|_{k-1} = F_k\, \hat{x}_{k-1}|_{k-1}, \quad P_k|_{k-1} = F_k\, P_{k-1}|_{k-1}\, F_k^T + Q_k \qquad (3)$$

where P is the state covariance and Q captures process uncertainty. EKF and UKF variants handle nonlinear motion or measurement functions.

For maneuvering UAVs, an IMM combines multiple kinematic models with mode probabilities; the classical IMM formulation remains attractive because it explicitly represents model switching while retaining bounded computation [31]. The important point is not that a neural network should replace these filters, but that any learned predictor should be judged against them under identical radar updates.

B. Data Association and Random Finite Sets

In clutter and multi-target scenes, the central problem becomes measurement-to-track association. Probabilistic data association soft-weights candidate measurements, while JPDA extends the principle to multiple targets [29]. MHT maintains multiple association histories and can preserve hypotheses through ambiguous crossings at greater computational cost [30].

Random finite set methods model target number and states jointly; PHD filters offer a computationally efficient intensity formulation, while labeled multi-Bernoulli/GLMB methods maintain identities more explicitly [32], [33]. OSPA provides a principled multi-object error metric that accounts for both localization and cardinality [34].

These methods are especially relevant to small UAVs because missed detections, clutter returns, and target crossings are not exceptional events; they are routine failure modes. A paper reporting only position RMSE on an already-associated single track does not demonstrate multi-target surveillance.

At minimum, multi-target evaluation should include false tracks, track fragmentation, identity switches, confirmation and deletion delay, and association behavior at crossings.

TABLE III. TRACKING FAMILIES FOR UAV SURVEILLANCE

Method	Strength	Failure / cost	Best role in cognitive system
EKF / UKF	Low cost; covariance is explicit	Model mismatch during abrupt maneuvers	Baseline estimator and uncertainty so urce
IMM	Handles multiple motion regimes	Model set and transition matrix must be tuned	Maneuver-aware baseline
PDA / JPDA	Probabilistic association in clutter	Crowded scenes increase combi natorics	Moderate multi-target ambiguity
MHT	Maintains alternative histories	Memory and pruning complexity	Difficult crossings / delayed evidence
PHD / GLMB	Principled birth/death/ cardinality	Implementation and parameter c omplexity	Dense or uncertain multi-object scen es
Hybrid learned tr acker	Can learn motion/context cues	Calibration and domain shift	Augment, not hide, estimator uncerta inty

V. LEARNING-BASED PREDICTION AND HYBRID ESTIMATION

A. From Classification to Trajectory Prediction

Learning has moved beyond target classification toward temporal reasoning over tracks. LSTM/GRU models can summarize recent kinematic history; transformers can model longer-range dependencies and interaction patterns; temporal convolutional networks can provide efficient sequence processing. Track-level methods are attractive because they use compact kinematic variables rather than large raw radar tensors, but they inherit any bias or fragmentation produced by the tracker.

Liu et al. [22] demonstrated the value of motion characteristics for radar bird/drone discrimination, while Tong et al. [28] combined temporal convolution, transformer modeling, and physics-constrained TimeGAN augmentation to address scarce and imbalanced trajectory examples.

The most important recent shift is toward joint tracking and prediction using high-dimensional radar. STIF integrates 4-D MIMO radar measurements, detection/tracking coupling, and an intention-aware transformer for micro-UAV trajectory prediction [35]. JTEA similarly integrates low-observable 4-D radar tracking with movement estimation rather than treating detection, tracking, and behavioral inference as unrelated post-processing stages [36]. These studies are significant because they move closer to the actual surveillance problem: weak observations must be accumulated across time before a reliable behavior estimate becomes possible.

B. Why Hybrid Estimation is Preferable

A learned model can be useful without becoming the sole estimator. Four hybrid patterns are particularly defensible: learning can (1) select among motion models, (2) adapt process noise Q, (3) generate a short-horizon trajectory prior, or (4) classify a maneuver that influences the next radar action. In all four cases the Bayesian filter remains responsible for covariance and measurement fusion. This architecture makes calibration visible and supports ablation: the learned component can be removed while the classical baseline still functions.

Pure sequence prediction often optimizes average displacement error without asking whether the predicted distribution is calibrated. For cognitive radar, calibration matters because a narrow but wrong predicted search sector can be worse than a broader conservative one. Prediction should therefore output a distribution or covariance-compatible uncertainty, and evaluation should include coverage probability or negative log-likelihood in addition to position error. A practical controller should expand the search region when the predictor is uncertain or out-of-distribution rather than treating every neural output as equally trustworthy.

VI. COGNITIVE RADAR AND RESOURCE MANAGEMENT

Cognitive radar converts uncertainty into action. The controllable action may include waveform, bandwidth, beam direction, dwell time, revisit interval, sensing mode, or task priority. In multi-function surveillance, time spent tracking known targets competes with time spent scanning for new ones. A generic cognitive policy can be written

$a_k = \pi(P_k, q_k, r_k, C_k, b_k)$ (4)

where P_k denotes tracking uncertainty, q_k track quality, r_k threat or mission risk, C_k scene context, and b_k the available time/energy budget. This formulation makes explicit that a cognitive action should depend on more than instantaneous tracking error.

Radar resource management has progressed from heuristics and information-driven scheduling toward RL and constrained RL. Charlish et al. [6] describe the transition from adaptive to cognitive RRM as a spectrum rather than a binary distinction. Thornton et al. [7] showed that deep RL can coordinate detection and tracking under spectral congestion. Durst and Brüggenwirth [8] applied deep RL to quality-of-service-based RRM. Transfer-based DRL has been explored for cognitive-radar task scheduling in dynamic environments [9]. Lu et al. [10] explicitly addressed the scanning-versus-multi-target-tracking trade-off; subsequent work treated the problem as multi-objective [11] and, in 2026, as constrained deep RL with an explicit time-budget constraint [12].

TABLE IV. EVOLUTION OF LEARNING-BASED RADAR RESOURCE MANAGEMENT

Study	Decision variable / setting	Contribution	Remaining issue
Thornton et al. 2020 [7]	Detection/tracking control in congestion	Deep-RL adaptive sensing under spectral competition	Transfer to different radar/ environment must be tested
Durst & Brüggenwirth 2021 [8]	QoS-based resource allocation	DRL reduces online optimization burden	Reward encodes engineering priorities
Akbar et al. 2024 [9]	Dynamic task scheduling	Transfer-based DRL improves adaptability	Task model may differ from UAV surveillance
Lu et al. 2024 [10]	Scan vs. multi-target track time	Learning-based joint allocation	Primarily algorithmic/ simulation evidence
Lu et al. 2025 [11]	Multi-objective time allocation	Pareto view of competing objectives	Operator selection of trade-off remains
Lu et al. 2026 [12]	Constrained track-while-scan budget	Explicit constraint handling with deep RL	Safety and robustness need field validation

A key lesson is that reward alone is not a safety constraint. In surveillance, a high long-term reward could still permit a locally unacceptable behavior such as starving the search task while servicing an uncertain track.

Constrained formulations are therefore more appropriate when scan coverage, total dwell, or latency must stay within hard bounds. Explainability is also becoming relevant: the operator should be able to tell whether extra radar time was allocated because covariance grew, the target maneuvered, the target approached a protected zone, or the policy encountered an unfamiliar state.

VII. AUTONOMOUS TARGET RE-ACQUISITION

Track loss is not a corner case for small-UAV radar. Temporary disappearance can result from occlusion, low SNR, a null in aspect-dependent scattering, multipath, CFAR thresholding, beam scheduling, or a target crossing. Classical track logic propagates the state during missed detections and deletes the track after a confidence or time threshold. That logic is reproducible but passive: it does not necessarily change how the radar observes the scene.

A cognitive system can make recovery an explicit closed-loop objective. When the detection is missed, the tracker propagates a state distribution; the predictor may supply a maneuver-conditioned prior; the controller converts the resulting uncertainty into one or more search sectors; and the radar increases revisit rate, dwell, or directional coverage until a candidate measurement is found.

The candidate then passes gating and identity checks before track continuity is restored. If confidence falls below a termination threshold or a resource budget is exhausted, the track is closed rather than being maintained indefinitely.

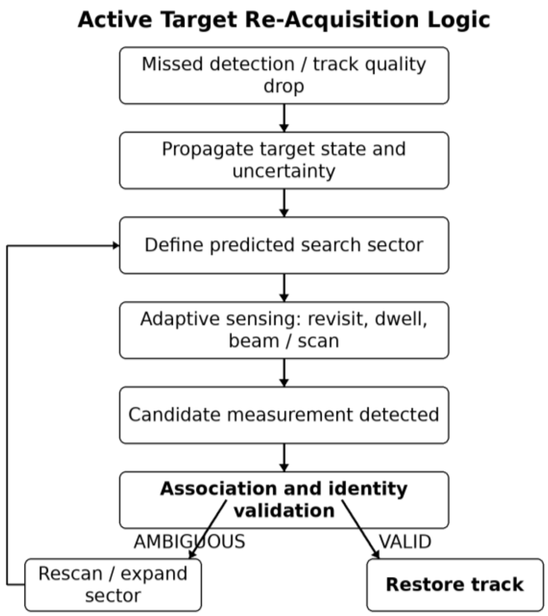

Fig. 2. Active re-acquisition logic. A missed detection increases uncertainty, which defines a predicted search sector and triggers an adaptive sensing action. Recovery is accepted only after association/identity validation; ambiguous candidates trigger another scan.

The literature is weaker on recovery metrics than on detection or classification. Many papers show a continuous sample trajectory but do not separately report how often a lost target is recovered, whether the identity is correct, or how long recovery takes. A recovery experiment should therefore include recovery probability, median and tail time-to-recovery, false recovery rate, wrong-ID rate, track fragmentation, and sensing resources spent during recovery.

Controlled occlusion and target-crossing tests are essential because an algorithm that recovers the wrong UAV quickly is not a successful tracker.

VIII. THREAT ASSESSMENT AND DECISION SUPPORT

Detection answers whether an object is present; classification estimates what it is; threat assessment asks what the observed behavior means relative to a protected asset and mission policy. Radar can contribute range, radial velocity, approach geometry, altitude, persistence, acceleration, maneuver pattern, and track confidence. These variables should be interpreted in context. A drone moving fast is not necessarily threatening if it is departing from the asset, while a slow persistent target near a restricted boundary may deserve attention.

Threat models in the recent literature include Bayesian networks for small-data conditions [40], threat-aware radar resource management [37], hierarchical fuzzy sensor fusion [38], and dynamic multi-asset geometric scoring with historical trend correction [39]. These methods are attractive when they expose contributing factors. A useful threat output is not only a discrete LOW/MEDIUM/HIGH label; it includes a calibrated score, dominant reasons, uncertainty, and the asset to which the score applies. This reduces the risk that classification confidence is incorrectly treated as threat probability.

For a radar-centered system, multimodal fusion should be viewed as optional evidence enrichment rather than a reason to hide radar limitations. EO/IR, RF, or acoustic confirmation can be valuable, but synchronization and calibration become new failure modes. The radar-only benchmark should therefore be reported first, followed by the improvement gained from fusion. Human oversight remains appropriate for security decisions because sensor data rarely encode the full operational context.

IX. INTEGRATED COGNITIVE SURVEILLANCE ARCHITECTURE

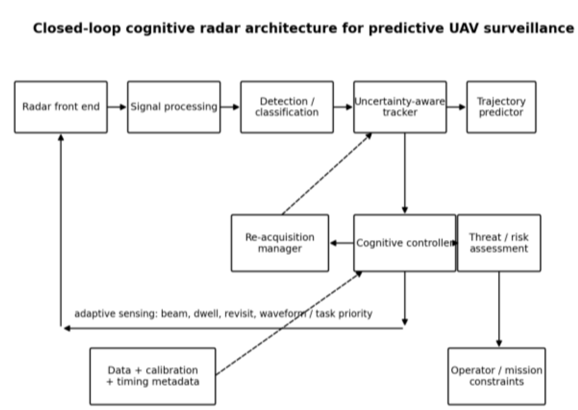

Fig. 3. Reference architecture for predictive cognitive UAV surveillance. The uncertainty-aware tracker and predictor feed a cognitive controller that adapts radar sensing. A separate recovery manager handles target loss, while threat assessment and mission constraints influence prioritization.

Figure 3 combines the strongest recurring ideas into one architecture. The radar front end supplies either raw/intermediate measurements or a target list. Signal processing creates range-Doppler/angle features, detections, and confidence indicators. Detection/classification proposes target hypotheses but does not own identity. Identity and uncertainty remain the responsibility of the tracker. The trajectory predictor operates on a history of filtered states and, where available, radar features. Its role is to enrich the motion model, not to erase covariance.

The cognitive controller receives tracking uncertainty, prediction, target priority, threat context, and resource constraints. It decides when and where to sense next. The re-acquisition manager is invoked when continuity degrades and requests a search action based on predicted state distributions. Threat assessment operates in parallel because a high-risk target may justify more frequent updates than a low-risk target at the same kinematic uncertainty. The architecture therefore couple's estimation and control while preserving modular evaluation: every block can be ablated and compared against a simpler baseline.

This separation also reduces a common source of overclaiming. A spectrogram classifier is a perception component, not a complete surveillance system. A trajectory transformer is a prediction component, not a proof of reliable tracking. A deep-RL scheduler is a cognitive-control component, not a proof of safety. The system claim becomes credible only when error and uncertainty are propagated across components and evaluated end to end.

X. DATASETS, BENCHMARKS, AND EVALUATION METRICS

A. Dataset Requirements

Dataset scarcity is one of the strongest cross-cutting limitations. Radar data are highly dependent on carrier frequency, waveform bandwidth, antenna aperture, mounting geometry, target aspect, clutter, and vendor preprocessing. Consequently, a model trained on one sensor or location may fail under another even when the UAV category is unchanged. The strongest benchmark design would release raw or minimally processed radar measurements together with derived maps/point clouds, synchronized ground truth, target metadata, and environmental metadata.

Ground truth should come from RTK/DGPS, calibrated optical motion capture, lidar, or another sufficiently accurate source. It should identify when a target is genuinely occluded or outside the sensor field of view. For multi-target experiments, target identities must be synchronized so that identity switches and false recovery can be measured.

If raw data cannot be released for security or licensing reasons, the paper should still document radar model, carrier band, bandwidth, chirp configuration, antenna geometry, update rate, detection thresholding, target types, flight speeds, altitudes, weather, and the train-validation-test split.

B. Multi-Layer Metric Set

TABLE V. MINIMUM MULTI-LAYER EVALUATION SET

Layer	Core metrics	What the metrics expose
Detection	P_d, P_fa, precision/recall, ROC, max/robust range	Sensitivity versus false alarms and range
Classification	F1, per-class recall, calibration, cross-domain accuracy	Class imbalance and generalization
Tracking	Position/velocity RMSE, OSPA, fragmentation, ID switches	Kinematic and association quality
Prediction	ADE/FDE, NLL/coverage, horizon-wise error	Accuracy plus uncertainty calibration
Recovery	Recovery rate, time-to-recovery, false/wrong-ID recovery	Performance after target loss
Threat	FPR/FNR, confusion by threat band, calibration, alert delay	Operational decision quality
Deployment	Update rate, end-to-end latency, RAM/flash, model size, power	Real-time and embedded feasibility

Reporting these layers together reveals trade-offs that isolated metrics hide. A resource policy may reduce track RMSE by dwelling longer on known UAVs while degrading new-target discovery. A classifier may achieve high accuracy on clean clips while failing after track fragmentation. A re-acquisition policy may increase recovery rate by widening the gate but also increase wrong-ID associations. End-to-end surveillance is therefore a multi-objective problem even before the cognitive controller is introduced.

XI. EMBEDDED AND REAL-TIME DEPLOYMENT CONSTRAINTS

The phrase real time should be tied to the radar update interval. End-to-end latency includes data transfer, FFT/beamforming or vendor preprocessing, detection, association, filter update, neural inference, cognitive decision, and display/command delay. If this sum approaches or exceeds the revisit interval, the controller is acting on stale information. Reporting only neural-network inference time is insufficient.

Embedded feasibility similarly requires more than model accuracy. Papers should report processor, precision, model parameters, memory footprint, batch size, and power where relevant. Quantization and pruning may reduce latency but can alter calibration, particularly in a predictor whose uncertainty directly determines a search sector. Edge-fog architectures can move heavy inference to a local computer, but communication delay and link failure become part of the system model. A robust cognitive radar should degrade gracefully: if the learned model or edge node is unavailable, a classical tracker and conservative scanning policy should continue operating.

This fallback principle is especially appropriate for M-Tech-scale prototypes. A low-cost radar module can provide kinematic target reports to an EKF/IMM and a lightweight predictor on an edge computer. The cognitive action can initially be a small set of interpretable scan/revisit modes rather than a high-dimensional waveform controller. Such a design still demonstrates cognition if the sensing action is explicitly driven by uncertainty and predicted target behavior.

XII. COMPARATIVE SYNTHESIS AND RESEARCH GAPS

The first major finding is that sensing and classification are more mature than integrated surveillance. There is now strong evidence that micro-Doppler, range-Doppler imagery, and learned radar representations can discriminate UAV-like targets in controlled and increasingly realistic scenarios [16]-[28]. However, these results do not automatically imply persistent identity or robustness to domain shift. The second finding is that tracking theory is mature but is often evaluated separately from modern learning. Bayesian filters provide exactly the covariance information needed by a cognitive controller, yet many deep predictors report deterministic errors without uncertainty calibration.

The third finding is that radar resource management is rapidly advancing. The progression from adaptive scheduling to deep RL, multi-objective RL, transfer learning, and constrained RL is clear [6]-[12]. The unresolved issue is external validity: most cognitive-control studies necessarily begin in simulation because the policy requires many trajectories and controlled counterfactual actions. For UAV surveillance, field measurements should eventually close this gap by replaying or live-testing the policy with real detection probabilities, clutter, beam mechanics, and sensor latency.

The fourth finding is that active re-acquisition is under-evaluated. Track-before-detect and integrated 4-D radar work [35], [36] show the benefit of using temporal evidence before hard thresholding, but recovery is still rarely a headline metric. The fifth finding is that threat assessment is becoming more dynamic and interpretable [37]-[40], yet threat labels are not standardized across UAV scenarios. The sixth finding is that deployment evidence remains weaker than algorithmic evidence, especially for memory, power, and end-to-end latency.

TABLE VI. RESEARCH GAPS AND TESTABLE RESEARCH ACTIONS

Gap	Why it matters	Recommended experiment
Sensor/domain shift	Radar signatures are hardware- and environment-dependent	Train on one site/UAV subset; test on unseen sensor setting, site, aspect, and UAV type

Gap	Why it matters	Recommended experiment
Uncalibrated prediction	Wrong narrow priors can misdirect sen sing	Report coverage/NLL; compare learned prior wit h IMM covariance and conservative fallback
Weak recovery evidence	Persistent surveillance depends on wha t happens after misses	Inject controlled 0.5-5 s target loss; measure reco very rate, time, wrong-ID and resource cost
Simulation-only cognitio n	Idealized dynamics hide false alarms a nd latency	Hardware-in-the-loop, recorded-radar replay, the n live radar with fixed and adaptive policies
Nonstandard threat label s	Cross-paper comparison is impossible	Publish scenario/asset definitions, factor weights, thresholds, calibration and operator rationale
Incomplete deployment profile	Offline accuracy can mask unusable lat ency	Report full pipeline latency, memory, CPU/GPU, sensor rate, communication delay and power

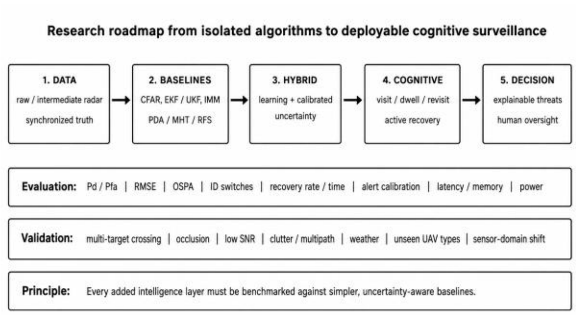

Fig. 4. Research roadmap for moving from isolated algorithms to a reproducible cognitive UAV-surveillance system. Each intelligence layer should be justified against simpler baselines and validated under failure-focused scenarios.

XIII. RESEARCH-QUESTION SYNTHESIS

A. RQ1 - Which sensing representations are most useful?

No single representation dominates all tasks. Range-Doppler maps are efficient for detection and provide a compact bridge from signal processing to CNN/transformer models. Micro-Doppler spectrograms are particularly informative for rotor-driven class discrimination, but are vulnerable to aspect and domain shift. Track sequences discard detailed scattering information but are computationally cheap and useful for behavioral classification. Four-dimensional MIMO radar retains the richest spatial-temporal structure among the reviewed compact-radar approaches, and recent STIF/JTEA results show that this structure can support low-observable tracking and early motion estimation [35], [36]. The practical choice should therefore follow the research claim: raw/intermediate data are necessary for novel signal processing; target lists are sufficient for tracking/resource-control research but cannot support claims about raw micro-Doppler processing.

B. RQ2 - Which tracking approach is the strongest baseline?

For a single maneuvering UAV, an IMM using constant-velocity, constant-acceleration, and coordinated-turn models is a stronger baseline than an EKF alone because it tests whether a proposed predictor truly adds value beyond multiple classical motion hypotheses. In multi-target scenes, the baseline must also model association uncertainty. JPDA is a reasonable middle ground for moderate target counts, MHT is valuable when ambiguous histories need to be preserved, and labeled RFS methods are appropriate when births, deaths, and identity must be modeled jointly [29]-[34]. The strongest paper is not the one with the most complex tracker, but the one that matches the tracker to the surveillance ambiguity and report's identity failures explicitly.

C. RQ3 - How should learning be integrated with prediction?

The evidence favors hybrid integration over unconditional replacement of the estimator. Learning is well suited to extracting latent motion modes, exploiting longer temporal context, and predicting short-horizon maneuver distributions. Classical filtering is well suited to measurement fusion and calibrated uncertainty. A publishable hybrid therefore exposes both components: the learned model proposes a distribution or model adaptation, and the filter updates it using measurements. This design supports confidence gating, out-of-distribution fallbacks, and interpretable ablation. STIF [35] is an important reference for end-to-end spatial-temporal integration, while track-level work such as [22], [28] illustrates the complementary value of behavioral features when raw radar data are unavailable.

D. RQ4 - What makes radar control genuinely cognitive?

Adaptivity becomes cognition when the sensing action is conditioned on an internal belief state and an objective, rather than on a fixed threshold alone. The belief state can include covariance, track quality, target behavior, threat, and remaining resource budget. Recent constrained-RL work is especially relevant because radar control is not an unconstrained game: minimum search coverage and maximum time budgets are engineering requirements [10]-[12]. For field systems, a learned controller should be compared against both a fixed schedule and a simple uncertainty-triggered heuristic. If RL cannot outperform the heuristic under realistic latency and false alarms, the additional complexity is not justified.

E. RQ5 - What should autonomous recovery optimize?

The recovery objective should be probability of correct identity restoration per unit sensing cost, not merely the first return of a detection. The predicted covariance should define the search geometry, while the cognitive controller selects revisit/dwell actions and the association logic rejects inconsistent candidates. Recovery evaluation should explicitly include crossings and distractors because these create the highest risk of fast but incorrect reacquisition. An effective controller should also know when to stop: indefinitely searching for a stale target wastes resources and can increase false associations.

F. RQ6 - What constitutes defensible threat assessment?

Threat assessment is defensible when it is asset-relative, time-varying, calibrated, and explainable. Useful factors include distance to protected assets, closing speed, approach direction, persistence, altitude, maneuvering, and track confidence. Bayesian and fuzzy approaches remain relevant because they expose relationships and uncertainty in small-data settings [38], [40]; recent dynamic multi-asset scoring further emphasizes temporal smoothing and approach geometry [39]. A learned threat classifier may be added, but the paper should still publish the label definition and show why the model raised an alert. Target class alone must not be equated with hostile intent.

TABLE VII. DIRECT ANSWERS TO THE SIX REVIEW QUESTIONS

RQ	Evidence-supported answer	Minimum validation
RQ1	Use the richest representation needed by the claim; 4-D/intermediate data aid joint perception, while target lists suit tracking/control.	Cross-aspect, cross-target and clutter tests
RQ2	IMM + association-aware multi-target method is a stronger baseline than isolated EKF tracking.	RMSE/OSPA + ID switches + fragmentation
RQ3	Hybrid learning-estimation preserves uncertainty while adding temporal context.	Classical baseline + calibration + ablation
RQ4	Cognition requires belief-driven sensing under explicit resource objectives/constraints.	Fixed and heuristic policies + real latency
RQ5	Recovery must optimize correct identity restoration, time, and sensing cost.	Controlled loss, distractors, wrong-ID rate
RQ6	Threat should be dynamic, asset-relative and explainable, not synonymous with class confidence.	Published labels, calibration, alert delay

XIV. REPRODUCIBILITY AND REPORTING CHECKLIST

A final IJIRT submission should make the experiment reproducible from the manuscript and supplementary files. The radar section should state carrier frequency, bandwidth, chirp/sweep configuration, antenna/MIMO geometry, field of view, range/velocity/angle resolution, mounting geometry, and update rate. The processing section should identify whether data are raw ADC, range-Doppler, point cloud, or target list and should state any vendor-side filtering. The tracking section should define state vector, process/measurement noise, confirmation/deletion logic, gates, and association method. The learning section should state dataset partitioning by flight/target, model size, optimizer, training seeds, early stopping, calibration method, and inference precision. The cognitive-control section should publish state variables, action space, reward/constraints, training environment, fallback policy, and resource budget.

Failure cases should be first-class results. At least one figure or table should summarize missed detections, false tracks, identity swaps, recovery failures, prediction failures, and alert errors. A clean trajectory is not evidence that the system is robust. Likewise, a mean latency value should be accompanied by worst-case or percentile latency because cognitive decisions can become stale during occasional compute spikes. Where a private dataset is unavoidable, authors should provide enough scenario metadata to permit a comparable reimplementation.

TABLE VIII. IJIRT-ORIENTED MINIMUM REPORTING ITEMS

Component	Required reporting items
Radar hardware	Band, bandwidth, waveform, antennas/MIMO, FoV, resolutions, update rate, mounting
Data	Representation, raw/vendor processing, target types, flight count, train/val/test split, ground truth
Tracking	State model, Q/R, gates, association, birth/confirm/delete logic, missed-detection handling
Learning	Architecture, parameters, input window, optimizer, seeds, calibration, inference precision/time
Cognitive control	State/action, reward and hard constraints, budget, baseline schedules, fallback behavior
Threat	Asset definition, features, label ontology, thresholds, explanation, human decision boundary
Deployment	Processor, RAM/flash, model size, communication delay, end-to-end latency, power where relevant

XV. IMPLICATIONS FOR LOW-COST 24-GHZ RESEARCH PLATFORMS

Many educational and prototype systems use compact 24-GHz FMCW modules because they are inexpensive and easy to interface with embedded processors. Their scientific value depends on what data the module exposes. If the device exports only range/velocity/angle target reports, the research should focus on tracking, trajectory prediction, uncertainty-driven revisit control, target-loss recovery, and decision support. It should not claim novel range-Doppler, CFAR, micro-Doppler, or raw waveform processing unless the underlying samples are actually accessible. This distinction protects the paper from a common mismatch between the stated algorithm and the sensor interface.

A strong low-cost experiment can still make a meaningful contribution. The radar target stream can feed an IMM or JPDA/GLMB tracker on a PC or edge device; a lightweight LSTM/TCN predictor can estimate short-horizon motion; and a cognitive controller can select among discrete sensing/scan modes using covariance and predicted behavior. Temporary loss can be emulated or measured by controlled occlusion, and recovery can be evaluated with the metrics in Table V. This architecture directly addresses the literature gap in active re-acquisition while remaining implementable without pretending that low-cost hardware provides capabilities it does not expose.

XVI. RECOMMENDED EXPERIMENTAL FRAMEWORK FOR FUTURE WORK

A. Baseline Stack

A rigorous experimental paper should begin with a deliberately simple baseline stack: CFAR or the sensor's fixed detector, an EKF/UKF, an IMM for maneuvering motion, and a fixed revisit schedule. Multi-target tests should add at least one association-aware method such as JPDA, MHT, PHD, or GLMB depending on computational scope. A learned predictor is then added as an ablation-controlled module. Cognitive control is evaluated last, comparing fixed scanning, uncertainty-triggered heuristic scheduling, and the learned policy under the same track histories.

B. Failure-Focused Scenario Matrix

Scenarios should be designed around failure causes rather than only nominal flights: low-SNR radial approach, tangential motion with weak Doppler, abrupt turns, altitude changes, temporary occlusion, UAVs crossing at similar range, clutter near buildings or vegetation, and an unseen target type. Weather or multipath variation should be included when the hardware supports safe outdoor testing. At least one scenario should deliberately force the system out of the predictor's training distribution to test fallback behavior.

C. Statistical Reporting

Results should be aggregated over repeated runs or trajectories and reported with variability. Mean error alone is inadequate when a small number of catastrophic track losses drive operational risk. Median, 90th/95th percentile, confidence intervals, and scenario-wise breakdowns are valuable. For learned models, hyperparameter tuning and test sets must be separated. For cognitive policies, evaluation seeds and scenario distributions should be disclosed because RL performance can be sensitive to training randomness and reward design.

D. Explainability and Safety Envelope

The cognitive policy should expose a safety envelope: minimum search coverage, maximum track dwell, maximum stale-track age, and conditions that force a return to a conservative scan. The threat module should expose contributing variables and should not autonomously infer hostile intent from target class alone. These design choices improve engineering traceability and align with the broader need for human-in-the-loop decision support in safety- and security-relevant surveillance.

XVII. THREATS TO VALIDITY

This review has four principal threats to validity. First, terminology is inconsistent: similar work may be described as cognitive sensing, adaptive radar, multifunction radar scheduling, track-before-detect, or intelligent resource management. Second, radar and defense literature are split across journals, conferences, technical reports, and restricted datasets, so no public search can guarantee complete coverage. Third, publication bias favors successful classifiers and policies, while negative field trials and failed generalization are less visible. Fourth, quantitative comparisons are limited by different sensors, targets, environments, and metrics; the review therefore avoids pooling accuracy numbers across incompatible datasets.

These limitations do not invalidate the synthesis, but they constrain what can be concluded. The strongest conclusions are structural: classical uncertainty-aware tracking remains essential; learned radar perception and prediction are increasingly effective but data-sensitive; cognitive RRM is technically mature enough for constrained, multi-objective policies; and end-to-end recovery, threat calibration, and embedded validation remain the least standardized parts of the pipeline.

XVIII. CONCLUSION

Cognitive radar provides a coherent framework for the next stage of small-UAV surveillance because it connects sensing to prediction and action. The reviewed evidence shows substantial progress in FMCW/mmWave perception, micro-Doppler recognition, Bayesian multi-target tracking, 4-D radar tracking and prediction, and learning-based radar resource management. Recent work through 2026 strengthens the case for hybrid systems: classical estimators supply calibrated state uncertainty, neural models learn complex motion or echo structure, and cognitive controllers allocate finite sensing resources according to uncertainty and mission objectives.

The central gap is no longer the absence of capable algorithms. It is the absence of standardized, failure-focused, end-to-end validation. Future work should treat re-acquisition as a primary metric, evaluate learned prediction with uncertainty calibration, compare cognitive policies against fixed and uncertainty-based baselines, publish sensor and ground-truth metadata, and report full pipeline latency and memory. Threat assessment should remain interpretable and context-aware. A system that satisfies these requirements would represent a credible transition from a collection of radar algorithms to a validated predictive cognitive surveillance architecture.

REFERENCES

[1] M. J. Page *et al.*, "The PRISMA 2020 statement: An updated guideline for reporting systematic reviews," *BMJ*, vol. 372, Art. no. n71, 2021, doi: 10.1136/bmj. n71.

[2] B. Kitchenham and S. Charters, "Guidelines for performing systematic literature reviews in software engineering," Keele Univ. and Durham Univ., Tech. Rep. EBSE-2007-01, 2007.

[3] S. Haykin, "Cognitive radar: A way of the future," IEEE Signal Processing Magazine, vol. 23, no. 1, pp. 30–40, Jan. 2006, doi: 10.1109/MSP.2006.1593335.

[4] S. Haykin, Y. Xue, and P. Setoodeh, "Cognitive radar: Step toward bridging the gap between neuroscience and engineering," Proceedings of the IEEE, vol. 100, no. 11, pp. 3102–3130, Nov. 2012.

[5] M. R. Bell, "Information theory and radar waveform design," IEEE Transactions on Information Theory, vol. 39, no. 5, pp. 1578–1597, Sep. 1993, doi: 10.1109/18.259642.

[6] A. Charlish, F. Hoffmann, C. Degen, and I. Schlangen, "The development from adaptive to cognitive radar resource management," IEEE Aerospace and Electronic Systems Magazine, vol. 35, no. 6, pp. 8–19, Jun. 2020, doi: 10.1109/MAES.2019.2957847.

[7] C. E. Thornton, M. A. Kozy, R. M. Buehrer, A. F. Martone, and K. D. Sherbondy, "Deep reinforcement learning control for radar detection and tracking in congested spectral environments," IEEE Transactions on Cognitive Communications and Networking, vol. 6, no. 4, pp. 1335–1349, Dec. 2020, doi: 10.1109/TCCN.2020.3019605.

[8] S. Durst and S. Brüggenwirth, "Quality of service-based radar resource management using deep reinforcement learning," in Proc. IEEE Radar Conf. (RadarConf21), 2021, doi: 10.1109/RadarConf2147009.2021.9455234.

[9] S. Akbar, R. S. Adve, Z. Ding, and P. W. Moo, "Transfer-based DRL for task scheduling in dynamic environments for cognitive radar," IEEE Transactions on Aerospace and Electronic Systems, vol. 60, no. 1, pp. 37–50, 2024.

[10] Z. Lu, M. C. Gursoy, C. K. Mohan, and P. K. Varshney, "Learning-based cognitive radar resource management for scanning and multi-target tracking," in Proc. IEEE Int. Conf. Commun. (ICC), 2024, pp. 2785–2790, doi: 10.1109/ICC51166.2024.10622225.

[11] Z. Lu, S. Kalia, M. C. Gursoy, C. K. Mohan, and P. K. Varshney, "Multi-objective reinforcement learning for cognitive radar resource management," in Proc. IEEE Int. Conf. Acoust., Speech Signal Process. (ICASSP), 2025, pp. 1–5.

[12] Z. Lu, M. C. Gursoy, C. K. Mohan, and P. K. Varshney, "Constrained deep reinforcement learning for cognitive radar resource management," IEEE Transactions on Radar Systems, vol. 4, pp. 627–644, 2026, doi: 10.1109/TRS.2026.3669518.

[13] V. C. Chen, F. Li, S.-S. Ho, and H. Wechsler, "Micro-Doppler effect in radar: Phenomenon, model, and simulation study," IEEE Transactions on Aerospace and Electronic Systems, vol. 42, no. 1, pp. 2–21, Jan. 2006.

[14] P. Molchanov, R. I. A. Harmanny, J. J. M. de Wit, K. Egiazarian, and J. Astola, "Classification of small UAVs and birds by micro-Doppler signatures," International Journal of Microwave and Wireless Technologies, vol. 6, nos. 3–4, pp. 435–444, 2014.

[15] S. Björklund, "Target detection and classification of small drones by boosting on radar micro-Doppler," in Proc. 15th Eur. Radar Conf. (EuRAD), 2018, pp. 182–185, doi: 10.23919/EuRAD.2018.8546569.

[16] A. G. Huizing, M. Heiligers, B. Dekker, J. de Wit, L. Cifola, and R. Harmanny, "Deep learning for classification of mini-UAVs using micro-Doppler spectrograms in cognitive radar," IEEE Aerospace and Electronic Systems Magazine, vol. 34, no. 11, pp. 46–56, Nov. 2019, doi: 10.1109/MAES.2019.2933972.

[17] B.-S. Oh, X. Guo, and Z. Lin, "A UAV classification system based on FMCW radar micro-Doppler signature analysis," Expert Systems with Applications, vol. 132, pp. 239–255, 2019, doi: 10.1016/j.eswa.2019.05.007.

[18] J. Park, D. H. Jung, K. B. Bae, and S.-O. Park, "Range-Doppler map improvement in FMCW radar for small moving drone detection using the stationary point concentration technique," IEEE Transactions on Microwave Theory and Techniques, vol. 68, no. 5, pp. 1858–1871, May 2020, doi: 10.1109/TMTT.2019.2961911.

[19] S. Rahman and D. A. Robertson, "Classification of drones and birds using convolutional neural networks applied to radar micro-Doppler spectrogram images," IET Radar, Sonar & Navigation, vol. 14, no. 5, pp. 653–661, 2020, doi: 10.1049/iet-rsn.2019.0493.

[20] M. Passafiume, S. Maddio, G. Collodi, and A. Cidronali, "Modeling small UAV micro-Doppler signature using millimeter-wave FMCW radar," Electronics, vol. 10, no. 6, Art. no. 747, 2021, doi: 10.3390/electronics10060747.

[21] R. M. Narayanan, B. Tsang, and R. Bharadwaj, "Classification and discrimination of birds and small drones using radar micro-Doppler spectrogram images," Signals, vol. 4, no. 2, pp. 337–358, 2023, doi: 10.3390/signals4020018.

[22] J. Liu, Q. Y. Xu, and W. S. Chen, "Classification of bird and drone targets based on motion characteristics and random forest model using surveillance radar data," IEEE Access, vol. 9, pp. 160135–160144, 2021, doi: 10.1109/ACCESS.2021.3130231.

[23] Y. Yang, F. Yang, L. Sun, T. Xiang, and P. Lv, "Echoformer: Transformer architecture based on radar echo characteristics for UAV detection," IEEE Sensors Journal, vol. 23, no. 8, pp. 8639–8653, Apr. 2023, doi: 10.1109/JSEN.2023.3254525.

[24] S. K. Han, J. H. Lee, and Y. H. Jung, "Convolutional neural network-based drone detection and classification using overlaid frequency-modulated continuous-wave range-Doppler images," Sensors, vol. 24, no. 17, Art. no. 5805, 2024, doi: 10.3390/s24175805.

[25] M. Larrat and C. Sales, "Classification of flying drones using millimeter-wave radar: Comparative analysis of algorithms under noisy conditions," Sensors, vol. 25, no. 3, Art. no. 721, 2025, doi: 10.3390/s25030721.

[26] P. Janpangngern et al., "High-resolution FMCW radar for small UAV detection using GNU software-defined radio," IEEE Access, vol. 13, pp. 86396–86412, 2025, doi: 10.1109/ACCESS.2025.3570635.

[27] T. Zhang and X. Song, "A classification algorithm of UAV and bird target based on L/K dual-band micro-Doppler and Mamba," Drones, vol. 10, no. 4, Art. no. 265, 2026, doi: 10.3390/drones10040265.

[28] F. Tong, K. Zhang, G. Liao, L. Li, J. Xu, and K. Jiang, "Radar UAV/bird trajectory feature classification based on TCN-Transformer and the PC-TimeGAN data augmentation framework," Sensors, vol. 26, no. 8, Art. no. 2528, 2026, doi: 10.3390/s26082528.

[29] Y. Bar-Shalom and E. Tse, "Tracking in a cluttered environment with probabilistic data association," Automatica, vol. 11, no. 5, pp. 451–460, 1975, doi: 10.1016/0005-1098(75)90021-7.

[30] D. B. Reid, "An algorithm for tracking multiple targets," IEEE Transactions on Automatic Control, vol. 24, no. 6, pp. 843–854, Dec. 1979, doi: 10.1109/TAC.1979.1102177.

[31] H. A. P. Blom and Y. Bar-Shalom, "The interacting multiple model algorithm for systems with Markovian switching coefficients," IEEE Transactions on Automatic Control, vol. 33, no. 8, pp. 780–783, Aug. 1988.

[32] R. P. S. Mahler, "Multitarget Bayes filtering via first-order multitarget moments," IEEE Transactions on Aerospace and Electronic Systems, vol. 39, no. 4, pp. 1152–1178, Oct. 2003, doi: 10.1109/TAES.2003.1261119.

[33] B.-T. Vo and B.-N. Vo, "Labeled random finite sets and multi-object conjugate priors," IEEE Transactions on Signal Processing, vol. 61, no. 13, pp. 3460–3475, Jul. 2013, doi: 10.1109/TSP.2013.2259822.

[34] D. Schuhmacher, B.-T. Vo, and B.-N. Vo, "A consistent metric for performance evaluation of multi-object filters," IEEE Transactions on Signal Processing, vol. 56, no. 8, pp. 3447–3457, Aug. 2008, doi: 10.1109/TSP.2008.920469.

[35] D. Huang, Z. Zhang, X. Fang, M. He, H. Lai, and B. Mi, "STIF: A spatial-temporal integrated framework for end-to-end micro-UAV trajectory tracking and prediction with 4-D MIMO radar," IEEE Internet of Things Journal, vol. 10, no. 21, pp. 18821–18836, 2023, doi: 10.1109/JIOT.2023.3244655.

[36] X. Fang, M. He, D. Huang, Z. Zhang, L. Ge, and G. Xiao, “JTEA: A joint trajectory tracking and estimation approach for low-observable micro-UAV monitoring with 4-D radar,” IEEE Transactions on Instrumentation and Measurement, vol. 73, pp. 1–14, 2024, doi: 10.1109/TIM.2023.3338650.

[37] M. I. Schöpe, H. Driessen, and A. Yarovoy, “Optimal threat-based radar resource management for multitarget joint tracking and classification,” Journal of Advances in Information Fusion, vol. 17, no. 1, pp. 46–62, 2022.

[38] F. Arapoglou, P. Zacharia, and M. Papoutsidakis, “Intelligent counter-UAV threat detection using hierarchical fuzzy decision-making and sensor fusion,” Sensors, vol. 25, no. 19, Art. no. 6091, 2025, doi: 10.3390/s25196091.

[39] Z. Zhang, L. Zhang, K. Huang, Z. Jia, M. Wu, L. Cui, and M. Zhang, “Three-dimensional dynamic UAV threat assessment using approach directionality and historical-trend correction for multi-asset protection,” Drones, vol. 10, no. 7, Art. no. 536, 2026, doi: 10.3390/drones10070536.

[40] R. Di, X. Gao, Z. Guo, and K. Wan, “A threat assessment method for unmanned aerial vehicle based on Bayesian networks under the condition of small data sets,” Mathematical Problems in Engineering, vol. 2018, Art. ID 8484358, 2018, doi: 10.1155/2018/8484358.