(Online) ISSN 2744-1741
Defense and Security Studies *Original Research*
Vol. 7, No. 2, 2026, pp.237-259
https://doi.org/10.37868/dss.v7.id346

Anti-jamming and secure communication for UAV swarms in contested environments: a systematic review

Mohammed H. Khalil[1*]
[1] Media Technology and Communication Engineering, University of Information Technology and Communications (UOITC), Baghdad, Iraq

*Corresponding author E-mail: mohammed.hussein@uoitc.edu.iq

Received: Jul 20., 2026
Revised: Sep 13., 2026
Accepted: Sep 14., 2026
Online: Sep 15., 2026

Abstract
The communication link that holds a UAV swarm together is one of the priority targets in modern electronic warfare, because without it the swarm loses the collective behaviour that makes it useful in the first place; when the link fails, coordinated action can break down within seconds. This systematic review looks at three main layers of defence that the research community has developed: waveform-level protection through Frequency Hopping Spread Spectrum (FHSS) and Direct Sequence Spread Spectrum (DSSS), channel-based protection through Physical Layer Security (PLS), and adaptive AI-based cognitive radio, including multi-agent reinforcement learning, which lets swarms detect and react to jamming with less need for human intervention. These three layers are not really competing with each other - each one deals with a different part of the threat (resisting jamming, resisting interception, and adapting to a changing adversary) - and this review concentrates on the gaps in deployment readiness and technical integration between them, working toward a communication security architecture for military UAV swarms that is closer to being usable in practice. The main finding of the review is a cautious one: every approach discussed here shows measurable gains in simulation, but the field has not yet produced a standardised, layered security architecture that both scales to swarm size and can hold up against an adversary that is itself AI-capable.

Keywords: UAV swarm communications; anti-jamming; frequency hopping spread spectrum; direct sequence spread spectrum; physical layer security; cognitive radio; multi-agent reinforcement learning; electronic warfare

1. Introduction

In practice, the communication link holding a UAV swarm together is the single most important point of failure for the whole system. Open-source reports on operations over eastern Ukraine since 2024 describe first-person-view (FPV) drone formations losing cohesion within seconds after entering contested airspace - not because they were shot down, but because Russian electronic-warfare (EW) systems cut the inter-platform links the formation needed to stay coordinated [5, 7]. Once isolated from each other, individual platforms fell back on pre-programmed behaviour, and in some reported cases began interfering with one another. These accounts come mostly from open-source defence reporting and released technical analyses rather than peer-reviewed

operational studies, and they should be read with that limitation in mind. They are broadly consistent, however, with the theoretical failure modes discussed in the peer-reviewed literature reviewed below, and this consistency is the reason they are cited here.

The communication link is more than a support function. It is what allows a swarm to achieve collective intelligence, synchronise action, share situational awareness, and combine platforms whose individual capability is limited into something that, as a group, is not [4]. From the point of view of an adversary, denying that link is also an efficient way to defeat a swarm, because a jamming platform is usually much cheaper to field than the swarm it disables.

Electronic warfare against UAV communications is not a new problem, but its character has changed considerably over the last decade. Low-cost jamming platforms are now widely available, and UAV swarms are being asked to operate at larger scale and in more demanding roles; these two trends reinforce each other. UAVs are now used for intelligence, surveillance and reconnaissance (ISR), precision strike, and logistics - functions with no easy substitute for reliable, real-time communication [30]. At the same time, adversaries are moving away from simple power-suppression jamming toward signal-aware attacks that can identify transmission patterns, track frequency hops in real time, and change their interference strategy faster than a static defence can respond [7].

In response, the research community has developed three broad layers of defence, which are largely complementary rather than competing. At the waveform level, Frequency Hopping Spread Spectrum (FHSS) and Direct Sequence Spread Spectrum (DSSS) improve resilience to jamming: FHSS reduces the effect of narrowband interference by switching the carrier frequency according to a shared pseudo-random sequence, while DSSS spreads each symbol over a wide bandwidth so that the jammer needs a higher processing gain to cause real damage [1, 3]. At the channel level, Physical Layer Security (PLS) treats the wireless channel itself as a security resource, exploiting the difference between the legitimate and the eavesdropping channel to limit interception rather than jamming. At the cognitive level, AI-based cognitive radio gives a UAV or swarm the ability to detect jamming, classify it, and select a countermeasure on its own [32]; Multi-Agent Reinforcement Learning (MARL) goes further, turning anti-jamming coordination into a behaviour that emerges from learning across the swarm rather than something engineered into a single platform [7, 21, 26].

This review treats these three layers not as competitors but as addressing different sides of the same problem - jamming resistance, interception resistance, and adaptive response — and asks where the gaps lie between what has already been demonstrated in the literature and what a deployable, hardened communication architecture for military UAV swarms would actually require. The conclusion reached here is a qualified one: each of the three layers shows measurable gains under the conditions in which it has been tested, but the field has not yet produced a validated, integrated security architecture that scales to swarm size and survives against an adaptive, AI-capable adversary. UAV swarm survivability is now a first-order consideration in planning across every operational domain - air, land, sea, space, and cyberspace - and it is this gap between demonstrated laboratory performance and deployable capability that this review is really about.

1.1. Review methodology

This review follows the PRISMA 2020 reporting guidelines for systematic reviews [36]. Records were collected from four databases - IEEE Xplore, Scopus, Google Scholar, and arXiv (cs.NI and eess.SP) - using combinations of the following search terms: UAV, unmanned aerial vehicle, drone swarm, anti-jamming, frequency hopping, FHSS, DSSS, spread spectrum, physical layer security, cognitive radio, deep reinforcement learning, multi-agent reinforcement learning, electronic warfare, and secure communication. The representative Boolean query used across the databases took the form ("UAV" OR "unmanned aerial vehicle" OR "drone swarm") AND ("anti-jamming" OR "physical layer security" OR "cognitive radio" OR "reinforcement learning") AND ("electronic warfare" OR "secure communication"), adapted to the field-search syntax of each database. Coverage was limited to records published between January 2017 and June 2026, with two exceptions: Wyner's wiretap-channel model and the original IEEE 802.11 spread-spectrum comparison, which were included outside this window because they form the technical basis of material discussed in Sections 4 and 3.

Eligibility was checked against four inclusion criteria at the full-text screening stage: (a) the study explicitly deals with UAV or drone communication systems, not wireless security in general; (b) it addresses at least one of the three security dimensions covered in this review - waveform-level anti-jamming, physical layer security, or AI-based cognitive radio/MARL; (c) the full text was accessible; and (d) for arXiv preprints specifically, the preprint had already been cited by at least one peer-reviewed publication, used here as a rough proxy for community vetting in the absence of formal peer review. Records were excluded at title/abstract screening if they dealt with wireless security issues unrelated to UAV or airborne platforms, or did not engage with any of the three dimensions above. The full screening and exclusion pathway is summarised in Table 1 and shown as a PRISMA 2020 flow diagram in Figure 1.

All screening, eligibility assessment, and quality scoring (Section 2) in this review were carried out by a single author. This is a genuine limitation for a systematic review: without a second, independent reviewer, inter-rater agreement could not be calculated, and the eligibility and quality judgements below should be read as one informed assessment rather than a consensus result. To reduce, though not remove, this subjectivity, the eligibility criteria above were fixed before full-text screening began, and every quality score reported in Section 2 was later re-derived in a separate, documented self-verification pass: each of the five rubric criteria in Table 2 was re-scored against a fixed set of scoring anchors, with a written justification recorded for every study (provided as supplementary material). This turns the quality-assessment scores from an unverified output into an auditable record of how each score was actually assigned, documenting the basis for every score reported below - though it does not substitute for genuine inter-rater verification, since both the original scoring and the verification pass were carried out by the same author. Independent, multi-reviewer verification remains a direction for a future update of this review. No formal risk-of-bias tool (e.g. a domain-specific adaptation of ROBIS) was applied to individual studies beyond the quality-assessment protocol in Section 2, which is noted here as a further limitation.

Table 1. PRISMA-style search and screening summary

Stage	Database / Source	Records	After Screening
Initial identification	IEEE Xplore	412	—
Initial identification	Scopus	287	—
Initial identification	Google Scholar	364	—
Initial identification	arXiv (cs.NI, eess.SP)	198	—
Total identified	All sources	1,261	—
Duplicates removed	—	—	1,089 unique records
Title/abstract screening (excluded: not UAV-specific or off-topic for the three security dimensions)	Relevance to UAV anti-jamming, PLS, or cognitive radio	—	347 retained
Full-text eligibility (excluded: not UAV-specific, did not address any of the three security dimensions, full text inaccessible, or not cited by peer-reviewed work for arXiv preprints)	Criteria (a)–(d), Section 1.1	—	37 eligible
Quality assessment (excluded: below quality threshold, Section 2)	Table 2 rubric	—	2 excluded (scores 3, 4)
Included in review	Peer-reviewed journals, IEEE Transactions, key conference proceedings, and vetted preprints	—	35 included

The rest of the paper is organised as follows. Section 2 describes the quality-assessment protocol applied to the studies retained after screening. Section 3 reviews FHSS and DSSS as waveform-level defences. Section 4 reviews Physical Layer Security, with particular attention to the line-of-sight-dominated UAV channel. Section 5 reviews AI-based cognitive radio and MARL approaches. Section 6 brings these three threads together, sets out a research-gap and technology-readiness analysis, and proposes a conceptual layered-architecture framework. Section 7 concludes with recommendations for the defence engineering community.

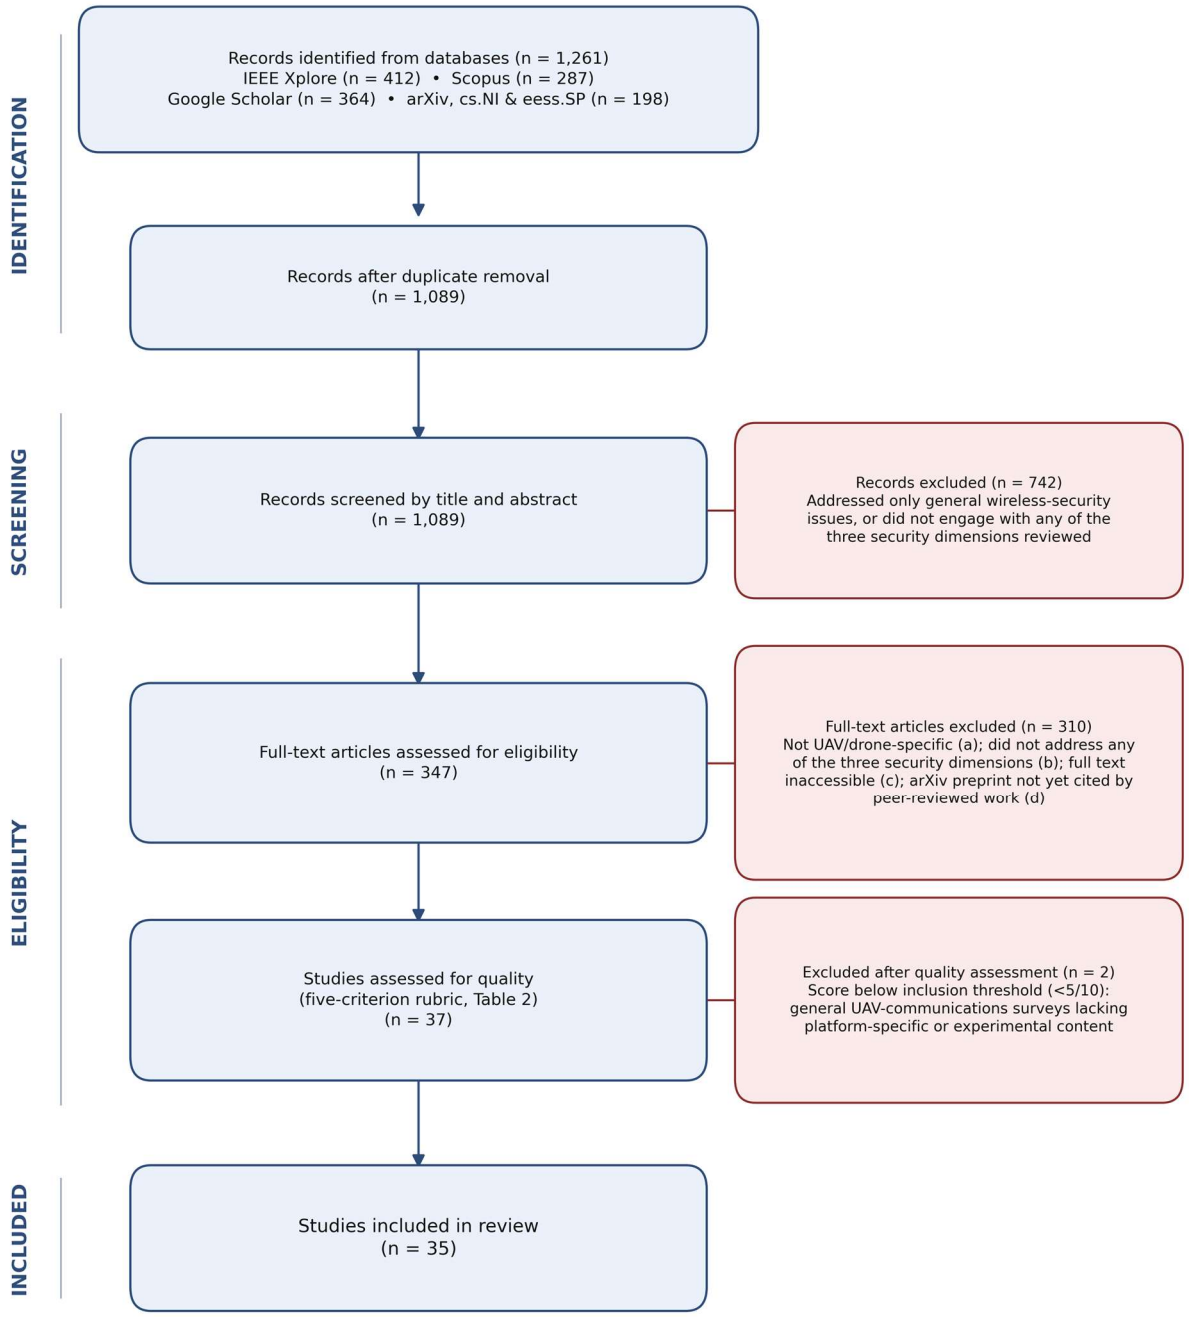

Figure 1. PRISMA 2020 flow diagram summarising record identification, screening, eligibility, and inclusion.

2. Quality assessment methodology

A quality-assessment protocol of this kind is necessary in a systematic review of engineering literature, because meeting the topical eligibility criteria alone (Section 1.1) does not tell us anything about how methodologically rigorous a study actually is. All 37 records that passed full-text eligibility screening (Table 1) were assessed one by one against the five-criterion rubric shown in Table 2, with each criterion scored 0–2, giving a maximum

possible score of 10. As mentioned in Section 1.1, the scoring was carried out by a single reviewer and later re-checked in a documented self-verification pass, with a written justification recorded criterion by criterion for every study. The scores reported below should therefore be understood as a transparent and auditable tool used to prioritise and organise this review, and not as a validated, reproducible quality metric of the kind a multi-reviewer Delphi-style protocol would produce.

Across the papers assessed, the mean overall score was 7.38 out of 10. Scores were highest on average for Novel Contribution (mean 2.00/2) and Quantitative Evaluation (mean 1.73/2), which suggests that the literature is generally original and reports quantitative results with care, at least within the scope each individual paper sets for itself. Scores were lower for Channel Realism (mean 1.19/2), and considerably lower for Experimental Validation (mean 0.78/2): only a few studies combine realistic propagation modelling with an actual Software-Defined Radio (SDR) implementation or hardware testing.

This pattern - strong on novelty and internal quantitative reporting, but weak on experimental and channel realism - reappears repeatedly through this review, and it is treated here as a structural finding rather than an incidental one. Two studies scored below the inclusion threshold used in this work (scores of 3 and 4) and were excluded from the final set of 35 reviewed papers; both were general UAV-communications surveys without platform-specific technical content or an experimental basis.

The two additional 2025-2026 references identified during revision to strengthen the coverage of recent literature (Section 5.6, Section 6.5; References [37]-[38]) were scored against the same rubric for consistency and are documented in the same supplementary verification record. Both meet the inclusion threshold (scores of 5 and 8, respectively), which is consistent with how they are treated in the main text, as directly relevant, quality-checked additions. They are reported here for transparency and not added to the count of 35 reviewed papers above, since they were identified after the original screening and quality-assessment pass described in Section 1.1 and Table 1. In this review they are cited as corroborating and illustrative recent evidence, rather than as part of the systematically screened corpus.

Beyond deciding what to include, the quality-assessment scores also shaped where this review places its critical attention. High-scoring papers clustered around Deep Reinforcement Learning (DRL)-based anti-jamming and Physical Layer Security (PLS) with trajectory optimisation, which is read here as an indication that these sub-areas are mature enough for a systematic comparison to actually be meaningful (Sections 4 and 5). The scores also point to a split in the literature: studies with the most rigorous channel modelling are rarely the same ones reporting SDR or hardware validation, and the two groups of studies cite one another only occasionally. Finally, no MARL study scored above 6, which reflects a literature that is technically inventive but has not yet combined that invention with channel realism or experimental validation; this point is returned to in the gap analysis in Section 6.7.

Table 2. Quality assessment framework for included studies

Criterion	Description	Score
Technical Relevance	Focus on UAV-specific anti-jamming, PLS, or cognitive radio in contested environments	0–2
Experimental Validation	Evidence of simulation, SDR implementation, hardware prototype, or field testing	0–2
Channel Realism	Realistic propagation model (Rician LoS, Rayleigh, Nakagami) and UAV mobility assumptions	0–2
Quantitative Evaluation	Clear reporting of BER, SINR, throughput, secrecy rate, or equivalent performance metrics	0–2
Novel Contribution	Presents a new algorithm, framework, or comprehensive comparative analysis beyond incremental extension	0–2

A caveat on comparing across studies follows directly from the propagation assumptions used in the reviewed literature. The 35 papers reviewed here differ in their assumed fading model (Rayleigh, Rician line-of-sight, Nakagami-m, or plain AWGN), jammer model, UAV mobility assumption, and network topology. Comparing reported gains in Signal-to-Interference-plus-Noise Ratio (SINR), Bit Error Rate (BER), or secrecy rate directly across studies that differ on several of these dimensions at once carries a risk of over-interpretation, since a technique that outperforms a given baseline under one set of modelling assumptions may not do so under another. The quantitative figures cited throughout this review, including those collected in Table 4, should therefore be read as results reported within each paper's own modelling context, not as directly comparable benchmark scores.

3. Waveform-level defence: FHSS and DSSS in contested UAV communications

FHSS and DSSS were developed decades before UAV swarms existed, and this fact tells us something useful. These waveforms are not used in modern UAV datalinks because they represent the most capable anti-jamming solution available; they are used because they are cheap to implement, they work reasonably well across a wide range of operating conditions, and - this matters for an engineering audience - their failure modes are well documented. Before discussing why the more advanced techniques covered in Sections 4 and 5 are actually needed, it is useful to understand where these failure modes come from.

3.1. The logic of each technique

FHSS works by evading the jammer in the time domain. The transmitter hops rapidly between narrowband channels according to a pseudo-random sequence known to the receiver [1]. If one frequency is jammed, communication continues on the remaining channels, and the constant hopping limits the amount of time interference on any single frequency can actually affect the link [2]. FHSS is relatively easy to implement, deals reasonably well with near-far interference, and is still the dominant choice in most commercial and military drone datalinks operating in the relevant bands. It is also convenient for swarm operation, since platforms sharing the same hopping sequence remain synchronised without having to renegotiate channel assignments continuously.

DSSS follows a different logic. Rather than moving across the spectrum, the signal stays on a single channel but is spread over a much wider bandwidth, which makes it difficult to detect and even more difficult to jam effectively. Each symbol is multiplied by a pseudo-random spreading code before transmission, and a receiver that knows this code can despread and recover the original signal, while a jammer without knowledge of the code has to overcome the full processing gain of the spreading operation before causing any real disruption [3]. Processing gain, defined as the ratio between the spread bandwidth and the information bandwidth, is directly related to jamming resistance: a larger spreading factor gives better jamming resistance but reduces the effective data rate [8]. This trade-off between throughput and resistance is a real design issue in swarm networks, which usually need to support several concurrent links carrying sensor data, video, and coordination traffic at once.

3.2. Where each technique fails

Used alone, neither FHSS nor DSSS is really suited to the contested environments that military UAV swarms are expected to operate in. Table 3 gives a comparison of the two techniques, while Table 4 collects the quantitative results reported in the literature reviewed in this section.

The main weakness of FHSS is reactive follower jamming. Even though the hopping sequence itself is not known to an adversary without the shared key, the sequence must still change fast enough that the adversary cannot detect the current frequency and redirect interference before the next hop occurs. State-level adversaries are now equipped with reactive jammers capable of detecting a transmission, identifying its frequency, and redirecting interference within microseconds [2], a timescale where raw reaction speed becomes more important than how random the hopping pattern actually is. The obvious countermeasure, increasing the hop rate, comes at the cost of higher hardware complexity, higher power consumption, and tighter synchronisation requirements across swarm nodes [2]. Underneath this problem there is also a more fundamental issue: FHSS depends on a shared secret key to generate and synchronise the hopping sequence, and once this key is compromised, whether

through cryptanalysis, physical capture of a platform, or weak key generation, the entire protection mechanism collapses [10].

DSSS fails for a different reason. When a high spreading factor is needed to keep the bit-error rate within acceptable limits under strong jamming, the processing gain that provides jamming resistance is paid for directly in terms of throughput. Wang et al. [8] report throughput reductions between 20% and 40% at high jammer-to-signal ratios, which can be a serious cost when the swarm is simultaneously carrying latency-sensitive coordination traffic. DSSS also relies on a shared spreading code in much the same way FHSS relies on a shared hopping sequence, and an adversary who manages to recover this code can correlate it with the spread signal and jam it coherently, achieving much higher efficiency compared with plain broadband jamming [10]. Both waveform techniques also share a limitation that has nothing to do with jamming resistance as such: neither one addresses interception. A signal that successfully resists jamming can still be received, decoded, and exploited by a passive adversary with sufficient signals-intelligence capability [10]. Closing this particular gap is the role of Physical Layer Security, which is reviewed in Section 4.

3.3. Hybrid approaches and operational realities

Neither FHSS nor DSSS has proven sufficient by itself in the electromagnetic environments that military UAVs now routinely operate in. The discussion of the Ukraine conflict presented below is based on open-source defence reporting and publicly released technical analyses rather than peer-reviewed operational data; full operational details remain classified, and the academic literature still lacks a consistent, independently verifiable evidence base for this case. It is included here mainly because it is widely cited across the research surveyed in this review, and because it is broadly consistent with the theoretical vulnerabilities described above, not because it constitutes an independently validated data point on its own.

With this caveat in mind, the Ukraine conflict remains the most widely discussed real-world case of FHSS-based UAV links operating against a state-level EW capability. Open-source reporting describes Russian EW systems degrading or blocking FHSS data links in the 900 MHz and 2.4 GHz bands, forcing Ukrainian operators into continuous adaptation [5]. This is consistent with, and is cited as motivation for, MARL research showing that swarms with coordinated, frequency-agile behaviour achieve considerably better communication survivability than swarms relying on static FHSS alone [7]. Ji et al. [11] combine null-space projection beamforming with adaptive frequency management, and report that this combination is more robust against the reactive jamming patterns described in operational reporting than either technique used separately.

One research direction that has emerged in response is hybridisation, combining the frequency agility of FHSS with the processing gain of DSSS [1]. Tao et al. [1] show that multi-slot, index-modulation-optimised FHSS can maintain a BER below 10^{-3} at 10 dB SNR under reactive jamming, although this comes at the cost of additional system complexity that is not negligible for resource-constrained UAV platforms. A different line of work addresses the predictability of the hopping sequence itself: de Curtò et al. [6] embed Quantum Random Number Generators (QRNGs) into the FHSS hop-sequence generator, which produces sequences that are not only pseudo-random but information-theoretically unpredictable, removing in principle the vulnerability to pattern learning. Whether QRNG-based FHSS can actually be implemented within the size, weight, and power (SWaP) limits of UAV platforms is still an open engineering question, but the existence of this line of work shows that reactive, learning-capable jamming is now treated by the research community as a first-order design threat rather than a marginal case.

Hybrid FHSS/DSSS approaches should be regarded as a floor for waveform-level defence rather than a ceiling. Both the operational reporting and the research literature agree that neither technique, alone or in combination, is sufficient to guarantee reliable swarm communication against an adversary able to observe, learn from, and exploit transmissions in real time.

3.4. Lessons learned from waveform-level defences

FHSS remains the most operationally mature anti-jamming technique available for UAV communications: its implementation complexity is low, its synchronisation requirements are well understood, and it has a long track

record of use in military datalinks. Three engineering conclusions follow from the reviewed literature, and each should be seen as a qualification of this maturity rather than a contradiction of it. First, the level of protection FHSS provides is limited by how quickly the adversary can react relative to the hop rate: against a jammer capable of sub-millisecond detection and response, closing this vulnerability window requires increasing hopping and synchronisation speed at the hardware level, not simply modifying the hopping algorithm. Second, QRNG-based hopping eliminates pattern-recognition vulnerability in principle, but it has not yet been validated against the SWaP constraints that govern most UAV platforms. Third, hybrid FHSS/DSSS architectures improve robustness against both reactive and broadband jamming, but they do not scale well, in terms of synchronisation overhead and bandwidth usage, as swarm size increases. These points should be treated as design requirements for the waveform layer of an integrated security architecture (Section 6.4), rather than simply as caveats attached to individual studies.

The conditions under which these results were obtained also deserve some attention. The strongest quantitative results reported in this section - BER below 10^{-3} at 10 dB SNR under reactive jamming [1], and 20–40% throughput loss under high jammer-to-signal ratio [8] - were obtained using simplified jammer models and single-adversary setups. Operational EW reporting from Ukraine, in contrast, describes multiple simultaneous jamming modes combined with adaptive power management and frequency tracking. At present, no validated benchmark exists in the research literature for a multi-modal, adaptive jammer operating at swarm scale. Until such a benchmark is established, the quantitative results reviewed in this section should be interpreted as performance envelopes obtained under idealised conditions, rather than as reliable predictions of operational performance.

4. Physical layer security: exploiting the channel against the eavesdropper

The techniques reviewed in Section 3 aim to keep communication available under active jamming. None of them answers a different question: what happens to a signal after it has already left the transmitter and reached a passive adversary? A receiver sensitive enough to intercept a jam-resistant signal can still decode and exploit it. Physical Layer Security (PLS) is the main technical answer the research community has proposed for this passive-interception problem.

4.1. The core principle

PLS is not a cryptographic layer added on top of the radio link. It treats security as something built into the channel itself, by exploiting the difference between the legitimate channel and the eavesdropper's channel. The foundational result, due to Wyner in the 1970s and extended considerably since then, is that when the legitimate channel is statistically better than the eavesdropping channel, information can be sent at a positive secrecy rate: a rate at which the legitimate receiver decodes correctly while the eavesdropper's received signal carries no usable information [14, 18, 19]. This is an information-theoretic guarantee, not a computational one - under the right channel conditions, no increase in an eavesdropper's computing power can change the outcome.

For UAV swarms, the appeal of PLS is structural as much as theoretical. Distributing and managing cryptographic keys across many independent, mobile platforms in a hostile environment is itself an attack surface: keys have to be generated, exchanged, stored, and revoked, and each of these steps can be compromised. A security mechanism based on channel physics rather than key secrecy removes this attack surface in principle. In practice, however, the channel conditions that PLS theory requires are often not available for UAV communications, for the reasons discussed next.

4.2. The line-of-sight problem

Terrestrial PLS analysis usually assumes that fading, shadowing, and multipath degrade the eavesdropper's channel more than the legitimate channel, and it is this difference that makes a positive secrecy rate possible. Low-altitude UAV operation over open terrain breaks this assumption in a structural way. The same strong, stable line-of-sight (LoS) propagation that makes reliable, long-range UAV communication possible also makes the transmitted signal easy to intercept from almost any direction [16]. In a broadcast configuration, the eavesdropper's received signal quality is not really any worse than that of the intended receiver, and interception

is possible from several kilometres away in open terrain, with no fading-based advantage available to the defender.

This is not simply an edge case for the PLS framework; it is close to the operating condition that UAV swarms present by default. Wang et al. [15] show that open operating geometry, dynamic node mobility, and broadcast-style swarm communication together create an interception vulnerability that conventional PLS analysis does not really capture. Platform-level constraints add to the problem: onboard computation, antenna array size, power budget, and payload capacity are all tightly limited on the small platforms that make large swarms economically viable, so a PLS technique that is theoretically sound but computationally or hardware demanding may not be practical on exactly the class of platform that needs it most.

4.3. Artificial noise and cooperative friendly jamming

Deliberate artificial-noise injection is the most developed PLS countermeasure for UAV systems: structured noise degrades reception at the eavesdropper without really affecting the intended link, since the interference is steered into the null of the legitimate receiver's channel [20]. Li et al. [20] show that generalised spatial directional modulation combined with artificial-noise precoding gives a meaningful secrecy-rate advantage in a UAV downlink scenario.

In a swarm, this technique gains a cooperative dimension: instead of requiring each transmitting node to generate its own noise, one dedicated platform can act as a friendly jammer on behalf of others that are carrying sensitive coordination traffic. Dang-Ngoc et al. [13] report an attainable secrecy rate of 3.1 bps/Hz for a cooperative friendly-jamming architecture in a swarm-assisted relay network, which is a substantial improvement over the near-zero secrecy rate achievable without PLS, and they combine this approach with wireless energy harvesting to offset the energy cost of cooperative jamming. It is interesting that the same LoS geometry that makes UAV links easy to intercept also helps the cooperative jammer, because LoS conditions make the eavesdropper's channel easier to estimate, which in turn improves the accuracy of the jamming direction [13].

One remaining dependency is that cooperative jamming usually still needs some knowledge of the eavesdropper's location or channel state. Nnamani et al. [12] address this point directly, and show through joint trajectory and power optimisation that secrecy rates of 1.2–2.4 bps/Hz are achievable even when the eavesdropper's position is completely unknown, results that are broadly comparable to scenarios where the eavesdropper's position is known exactly. Solving the joint trajectory-and-power optimisation problem in real time, across a mobile swarm operating under adversarial conditions with limited onboard compute, remains an open problem that the existing literature does not fully address. Wu and Zhang [17] extend this analysis to a two-UAV configuration and confirm that meaningful secrecy rates are theoretically achievable with trajectory optimisation, but they also flag a scalability concern, since computational complexity grows quadratically with the number of UAV nodes.

4.4. Beamforming, IRS, and emerging directions

Two further PLS directions go beyond artificial noise. The first, beamforming, concentrates transmitted energy toward the intended receiver while reducing leakage in other directions, which improves both anti-jamming and anti-interception performance at the same time [20]. Effective beamforming needs multiple antenna elements, and the size, weight, and aerodynamic constraints of UAV platforms limit the array size accordingly; multi-antenna beamforming stays more attractive in principle than in practice for the smaller, lower-cost platforms that make up most UAV swarm architectures.

The second direction, Intelligent Reflecting Surfaces (IRS), uses passive, programmable reflective elements to reshape the wireless channel in favour of the legitimate link and against the eavesdropper. IRS can be jointly optimised with UAV placement to improve the legitimate channel while degrading the eavesdropper's channel [22]. Xue et al. [22] report secrecy-rate improvements of 15–30% over UAV-only PLS schemes in an IRS-assisted configuration, at a noticeably lower power cost than active transmission-based methods. The limiting factor here is infrastructure: IRS elements have to be physically deployed in advance, which assumes some

degree of pre-established infrastructure or environmental stability that a fast-moving, contested battlefield may not offer, and the literature has not really addressed this deployment problem in a direct way.

Looking further ahead, 6G and integrated sensing-communication-computing-control (ISAC) architectures represent a longer-horizon direction for UAV PLS. Khan et al. [21] argue that AI-based channel estimation and ultra-reliable, low-latency communication primitives in future 6G systems could change the channel-quality difference that PLS requires, and this could potentially loosen the LoS constraint described in Section 4.2. This is still a research direction rather than a deployed capability, but it suggests that the LoS constraint may not be permanent.

4.5. What the PLS literature has not resolved

Most published PLS work on UAVs, including several of the studies reviewed above, explicitly limits itself to a single UAV or a small, fixed multi-node configuration, and states this as a limitation rather than treating it as incidental. Existing frameworks do not fully capture the dynamics of a large, distributed, continuously repositioning swarm, where changing node roles and formation add coordination overhead and channel variability that single-node or static-configuration analyses do not model. Where the eavesdropper is itself mobile and deliberately tracks the swarm to maintain a good intercept, the joint power-allocation and trajectory-control problem becomes much harder than any single-node analysis suggests [15]. Underlying all of this, and rarely addressed directly in the PLS literature, is a systems-integration problem: PLS mechanisms have to share power, bandwidth, and computation with the spread-spectrum waveforms and cognitive radio systems discussed in Sections 3 and 5, and this resource competition is an architectural issue that a security system needs to address explicitly rather than absorb as a side effect.

4.6. Lessons learned from physical layer security

Three engineering conclusions follow from the PLS literature reviewed here. First, the LoS-dominated propagation geometry of low-altitude UAV operation is not a minor implementation detail but a structural obstacle to the channel-quality difference that PLS theory requires; it should be regarded as the central unsolved problem in UAV PLS rather than an assumption that can simply be relaxed. Second, the reported secrecy-rate gains are real but limited: cooperative friendly jamming reaches up to 3.1 bps/Hz against a near-zero baseline without PLS [13], and trajectory optimisation against an eavesdropper of unknown position achieves 1.2–2.4 bps/Hz [12]; these figures should be read as achievable upper bounds under the stated assumptions, and they are likely to get worse as swarm size and node mobility increase. Third, IRS improves on UAV-only PLS by 15–30% [22], but it is better understood as suited to semi-permanent or pre-configured environments than as a general solution for contested, infrastructure-poor deployments. Taken together, these three findings suggest that PLS is best positioned as one component of a multi-technique security architecture, with its resource requirements negotiated explicitly against the other layers reviewed in this paper, rather than as a stand-alone solution.

5. AI-Based cognitive radio: from reactive switching to autonomous anti-jamming

The waveform-level defences and the PLS techniques discussed in Sections 3 and 4 are, in essence, static. They are configured before deployment against an expected threat and then used as they are. If an adversary manages to learn the hopping pattern, adapt to the beamforming null, or otherwise defeat the fixed configuration, neither of these layers has a built-in mechanism to react. AI-based cognitive radio is the response that the research community has proposed to this problem.

5.1. Cognitive radio: system overview

Cognitive radio spans a wide range of sophistication, starting from simple energy detection with threshold-based channel switching and going up to advanced systems that can characterise jammer behaviour, predict its next move, and select a countermeasure within milliseconds. In this review, the term cognitive radio is used for systems at the more capable end of this range, i.e., platforms that sense the electromagnetic environment, reason about what they observe, and take action without human intervention. While FHSS and DSSS remain fixed against an adversary that keeps adapting, a cognitive-radio-enabled UAV continuously senses the spectrum,

updates an internal model of the interference it observes, and decides in real time which frequency, power, waveform, timing, and coding rate should be used. Nawaz and Alzahrani [34] report a jamming classification accuracy of 97.3% at SNR ≥ 5 dB using machine-learning-based cyclostationary spectral analysis, while Imran et al. [27] show that particle-swarm-optimised deep learning can achieve competitive jamming-detection performance with lower computational cost, which is an important point for resource-constrained UAV platforms.

5.2. Deep reinforcement learning as the primary engine

In most recent work, the anti-jamming decision problem is formulated as a Markov Decision Process (MDP), where a UAV agent observes the electromagnetic environment, selects a transmission parameter, and receives a reward that is usually based on throughput, BER, or link reliability. Deep Reinforcement Learning (DRL) goes beyond the classical tabular Q-learning by approximating the value function with a deep neural network, and this allows some generalisation to channel and jamming conditions that were not explicitly seen during training [23]. Feng et al. [23] report SINR improvements of 5–12 dB and throughput gains of 30–60% compared with static frequency-assignment baselines. Yang [28] reports that DRL performs better than heuristic and classical adaptive methods for most jammer types, including reactive jammers. Xiao et al. [26] extend DRL-based anti-jamming to the application layer by training a policy that preserves video-streaming quality of experience under jamming, rather than optimising only link-layer metrics, and Dong et al. [29] apply a similar idea at millimetre-wave frequencies with joint beamforming and power control, indicating that the DRL approach is not limited to sub-6 GHz channels.

Designing the reward function is a real engineering decision that has operational consequences. A reward function that maximises throughput alone will train agents that give low priority to link reliability whenever jamming is mild, and such an agent may then perform poorly once the jamming severity exceeds what it experienced during training. More robust formulations combine several objectives together — throughput, link continuity, energy consumption, and inter-node interference - as discussed in [25]. Another approach relies on deception instead of pure evasion: Van Huynh et al. [9] train a deep dueling network to broadcast decoy traffic on channels that, from the point of view of a monitoring jammer, seem to follow a predictable pattern, so that the jammer's attention is drawn away from the channels that actually carry coordination traffic. In general, this line of work treats the adversary as a strategic actor that can be manipulated, not simply avoided, which is a fundamentally different posture compared with the static waveform designs reviewed in Section 3.

5.3. Multi-agent coordination: the swarm dimension

Single-UAV cognitive radio does not provide something that becomes essential at swarm scale, which is coordinated anti-jamming response across multiple agents. A single agent can only observe the channels visible from its own position, and this does not necessarily represent the electromagnetic environment as a whole. A swarm of agents that share observations and coordinate channel assignment can, on the other hand, build a distributed picture that no single agent could obtain by itself. Formalising this kind of collective intelligence - shared learning experience, jointly evaluated channel-selection strategies, and a collectively executed anti-interference policy - is an active research area [7, 26]. Zeng et al. [35] show that federated learning allows UAVs to update a shared policy without centralising the training data, reducing the training-phase communication overhead by 40–60% while keeping the policy quality almost the same, which is directly relevant for deployments where training-time bandwidth is itself limited.

Centralised Training with Decentralised Execution (CTDE) is the dominant architecture in this literature. A centralised critic uses the global state, observations, actions, and rewards of all agents during training, while at execution time each agent acts based on its own local observation, without needing to stay in contact with a central controller. This avoids the single point of failure that a jamming-vulnerable central coordinator would represent. Peng et al. [4] evaluate CTDE-trained multi-agent frameworks for collaborative relay configurations with swarm sizes between 4 and 12 agents, and report a swarm throughput roughly 40% higher than single-agent baselines. More advanced MARL frameworks treat transmission parameters, power allocation, and UAV trajectory as one joint optimisation problem [25]. Yin et al. [25] show that, for small swarms, jointly optimising

communication and control gives better results than optimising the two separately, both in terms of communication performance and energy efficiency. Khan et al. [21] place these results within a longer-term ISAC/6G perspective, arguing that future infrastructure could provide the low-latency coordination primitives that swarm-scale MARL requires.

5.4. The adversarial arms race

A parallel body of literature models the jammer-UAV interaction as a game between rational agents instead of a learning problem. Game-theoretic formulations give an explicit model of adversary decision-making and allow equilibrium analysis. Stackelberg (leader–follower) formulations are the most commonly used approach in this context [24], and they are typically applied to determine how much advantage a jammer gains by observing and learning the UAV's transmission policy before committing to its own strategy, and to identify design choices that can reduce that advantage.

One limitation of purely game-theoretic approaches is that they require the utility functions and strategy spaces to be defined in advance, which is not easy when the adversary's goals and capabilities are not fully known. Hybrid approaches try to deal with this by combining a structured game-theoretic model of the adversarial interaction with reinforcement learning that updates the model based on observed experience [24]. Krayani et al. [24] show that a cognitive UAV using active-inference-based resource allocation, based on the free-energy principle from computational neuroscience, can outperform both DRL and game-theoretic baselines in rapidly changing jamming environments. Krayani et al. [18] also introduce complementary self-awareness models for joint physical-layer security and cognitive anti-jamming.

5.5. Where the approach still needs to improve

Simulation results in the cognitive-radio and MARL literature are consistently strong, but there are three gaps that separate the demonstrated laboratory performance from actual field deployment.

The first and most basic gap is sensing reliability. Every cognitive-radio anti-jamming system depends on correctly characterising the type, power, and spatial location of the jamming signal before it can make any decision at all. Nawaz and Alzahrani [34] report 97.3% classification accuracy at SNR $\geq$ 5 dB, but the accuracy drops sharply below this threshold, and contested real-world environments often operate at SNR values below 5 dB. UAV platforms make this problem worse because of the limited sensor-payload bandwidth, constrained onboard computing power, and Doppler effects that sensing algorithms designed for stationary platforms do not take into account. A cognitive-radio system that misclassifies noise as jamming, or that fails to detect a low-power reactive jammer, will make systematically wrong anti-jamming decisions, no matter how well its downstream policy network has been trained.

The second gap concerns the adversarial feedback loop. A cognitively capable jammer can observe a UAV's transmission behaviour, infer its policy, and adapt its own jamming strategy to exploit any regularity in that policy [18]. More seriously, a sophisticated adversary could deliberately poison the feedback that a learning agent receives during operation, pushing it towards a policy that appears to perform well against the poisoned feedback but that actually performs poorly against the adversary's real strategy. This adversarial-machine-learning attack surface has not been studied in a systematic way in the UAV anti-jamming literature, which means the published robustness claims may be more optimistic than they should be when facing an adversary that is specifically trained to exploit the learning process.

The third gap is scale. Most published MARL results use swarms of 4-12 agents evaluated against simplified channel models, while large and dynamic swarms introduce intra-swarm interference, resource contention, and coordination overhead that grow non-linearly with the swarm size [25]. Generalising reliably from small-swarm training to large-swarm operation is an open problem in multi-agent learning in general, and it becomes even harder here because the learning task is coupled to a physical channel whose complexity also grows with the number of active nodes. None of the studies reviewed here demonstrate a credible path from a twelve-agent simulation result to a sustained, coordinated anti-jamming performance across a hundred-agent swarm under continuous electronic attack.

5.6. Lessons learned from AI-driven cognitive radio and MARL

AI-based cognitive radio and MARL represent the most technically advanced anti-jamming research reviewed in this paper, and the simulation results obtained are genuinely strong. However, a consistent pattern can still be seen behind these results: strong performance is demonstrated against the jammer type, channel conditions, and sensing assumptions that the system was trained under. The three gaps identified in Section 5.5 - unreliable sensing below 5 dB SNR, vulnerability to adversarial manipulation of the learning process, and the lack of validated results beyond roughly a dozen agents - mark the boundary between what DRL and MARL have actually proven and what an operational anti-jamming system would actually need.

Three engineering conclusions can be drawn from this discussion. First, DRL reliably improves SINR by 5–12 dB and throughput by 30–60% in simulation, and this has been confirmed by several independent studies [23, 28], but these results assume simplified, largely stationary jammer models and near-perfect spectrum sensing; how DRL compares with FHSS under a genuinely adaptive, sensing-aware adversary has not been rigorously quantified yet. Second, CTDE appears to be the most promising architecture for swarm-level anti-jamming, since it removes the single-point-of-failure dependency on a central coordinator during execution, but its demonstrated throughput advantage has so far only been shown for swarms of up to roughly twelve UAVs, and it is not known whether this advantage still holds at operationally relevant swarm sizes. Third, the adversarial-manipulation attack surface, where the reward signal of a defending agent is manipulated in order to induce an exploitable policy, is a credible but largely unstudied threat, and the published robustness claims should be treated with some caution until adversarial co-training against a learning opponent becomes a standard part of the evaluation pipeline for AI-based anti-jamming systems.

A recent and comprehensive survey on agent-based UAV anti-jamming techniques, published after most of the literature reviewed above, reaches a similar structural conclusion independently: it organises the field around a perception-decision-action loop and identifies robust, low-SNR spectrum sensing and coordinated, hybrid-attack-resilient defence architectures as open problems rather than solved ones [37]. The fact that a 2025 survey converges on the same sensing and cross-layer gaps that were identified independently in Sections 5.5 and 6.7 of this review is considered here as corroborating evidence that these gaps are structural to the field, rather than being artefacts of the particular sample of studies chosen for this review.

6. Toward a layered security architecture for UAV swarms

If Sections 3–5 are read one after another, they can give the impression of increasing sophistication, as if each layer solves what the previous layer could not, with cognitive radio appearing as a final, culminating solution. The evidence reviewed here does not really support this reading. FHSS/DSSS, PLS, and AI-driven cognitive radio are not successive generations of the same technology. They address different threat dimensions, operate on different timescales, and rely on different platform resources. Table 3 sets out this structural comparison, and Table 4 collects representative quantitative evidence from the reviewed literature.

Table 3. Comparative analysis of anti-jamming and secure communication techniques

Technique	Jamming Resist.	Intercept Resist.	Complexity	Energy Cost	Scalability	Maturity
FHSS	High (static)	Low	Low	Low	High	High (deployed)
DSSS	Moderate-High	Low	Low–Moderate	Moderate	Moderate	High (deployed)
Hybrid FHSS/DSSS	High	Low	Moderate–High	Moderate–High	Moderate	Moderate
PLS - Artificial Noise	Low-Moderate	High	Moderate	High	Low	Moderate (research)

Technique	Jamming Resist.	Intercept Resist.	Complexity	Energy Cost	Scalability	Maturity
PLS - Beamforming	Moderate	High	High	High	Low	Low–Moderate
PLS - IRS	Moderate	High	High	Low (passive)	Low	Low (emerging)
Cognitive Radio (DRL)	High (adaptive)	Moderate	High	High	Moderate	Low–Moderate
MARL Swarm	Very High	Moderate–High	Very High	High	Low (current)	Low (simulation)

Table 4. Representative quantitative metrics from reviewed literature

Technique / Study	SINR / BER	Throughput	Secrecy Rate	Detect. Acc.	Energy
FHSS (Tao et al. [1])	BER < 10^{-3} at SNR 10 dB under reactive jamming	~85% of nominal	N/A	N/A	Low overhead
DSSS (Wang et al. [8])	3–8 dB SINR gain over unjammed baseline	Reduced 20–40% at high JSR	N/A	N/A	Moderate
PLS Art. Noise (Nnamani et al. [12])	N/A	N/A	1.2–2.4 bps/Hz secrecy rate	N/A	High (trajectory opt.)
PLS Coop. Jamming (Dang-Ngoc et al. [13])	N/A	N/A	Up to 3.1 bps/Hz vs 0 without PLS	N/A	Moderate (energy harvest.)
DRL Anti-Jamming (Feng et al. [23])	SINR improvement 5–12 dB vs static	Throughput 30–60% higher than FHSS	N/A	Jammer detect. ~92%	High (inference cost)
MARL CTDE (Peng et al. [4])	N/A	Swarm throughput 40% > single-agent	N/A	N/A	Scalable up to 12 agents
Spectrum Sensing (Nawaz & Alzahrani [34])	N/A	N/A	N/A	Classification acc. 97.3% at SNR ≥ 5 dB	Low (cyclostat. analysis)

6.1. Different threats, different tools

The clearest qualification to any comparative reading of this literature is that the three approaches counter different threats, so combining them without an explicit architecture leaves predictable blind spots. FHSS and DSSS keep the communication link working under active jamming. PLS is designed mainly to prevent interception, denying an adversary the ability to extract useful information from a signal it manages to receive. Cognitive radio is the only one of the three built to work against both threat classes at once, because its defining capability - detecting, classifying, and reacting to an adversary whose behaviour is not fully known beforehand - is threat-agnostic by design.

None of this establishes a technical ranking among the three techniques. A UAV swarm operating in a contested EW environment faces all three threats simultaneously: jammers target the inter-platform links, signals-intelligence units attempt to extract operational value from intercepted signals, and sophisticated EW systems

adapt their tactics as an engagement develops [30]. A swarm that relies on waveform-level defence alone can resist jamming but remains exposed to a patient eavesdropper with a wideband receiver. A swarm relying on PLS alone can resist interception, yet it can still be jammed by a competent broadband jammer. A swarm relying on cognitive radio alone gains adaptivity, but it also inherits cognitive radio's own dependability problems, mainly that its decisions are only as good as the accuracy of its sensing and classification.

6.2. The conflict between security and resource consumption

All three security layers draw on the same limited pool of UAV swarm resources: onboard computation, transmission power, available bandwidth, and decision latency. These budgets are not independent, so allocating resources to one layer necessarily takes them away from the others.

The clearest tension is between PLS and cognitive radio. Artificial-noise injection consumes transmission power that a cognitive radio system needs to maintain adequate SINR on the channels it is monitoring. PLS beamforming uses antenna and processing resources that could otherwise support the rapid frequency changes required by cognitive radio's waveform-agility function. The spectrum-sensing pipeline of a cognitive radio also has to continuously track and compensate for UAV movement caused by trajectory optimisation performed for secrecy-rate maximisation [31]. These interactions are not simply incidental friction: the performance of a combined PLS-plus-cognitive-radio system cannot be predicted by adding up the performance of each component measured separately, and none of the frameworks in the reviewed literature treats PLS and cognitive radio as a jointly designed and jointly evaluated system.

A related tension appears at the waveform layer. DSSS's processing gain does not come for free; obtaining reliable jamming classification requires wideband channel monitoring, and the cost of this rises with jammer power. FHSS's frequency agility can also conflict directly with the channel dwell time that spectrum sensing needs to build a reliable estimate of the interference environment. An agent that hops before its sensing algorithm has collected enough samples is forced to make anti-jamming decisions based on incomplete information. Coordinating the timing between waveform-level hopping and cognitive-radio sensing matters, but the reviewed literature has not treated it as a joint design problem.

6.3. The simulation-to-reality gap

A pattern recurring across the studies reviewed is strong performance under the conditions a system was designed for, together with comparatively little evidence of performance under conditions it was not designed for. This is not a criticism of any individual paper. It reflects a structural feature of how the field has developed, shaped by the practical need to evaluate against tractable, well-defined baselines rather than against the full complexity of an operational environment.

The gap is widest in the cognitive-radio literature, where adversary models are generally simpler and more static than the adversaries that originally motivate this line of research. It also appears in the PLS literature, where channel models often assume perfect channel-state knowledge, or draw estimation errors from a simplified statistical distribution that does not capture the changing conditions a mobile swarm actually experiences in complex terrain. It appears again in hybrid FHSS/DSSS work, where jamming-detection latency, synchronisation overhead, and multi-platform interference are frequently abstracted away. Closing this gap will need more than higher-fidelity simulation. It needs standardised benchmarks in which an adaptive adversary, imperfect sensing, platform resource limits, and swarm-scale coordination overhead are treated as first-class experimental variables rather than simplifying assumptions.

6.4. A conceptual framework for a layered security architecture

The discussion above points to what an integrated, multi-layer security architecture for a military UAV swarm would need to look like. The inter-layer integration requirements are mostly implicit in a literature that typically studies each security layer in isolation. Table 5 makes these requirements explicit, classifying each layer by its constituent techniques and the threat it mainly addresses. It should be made clear what follows here is: this is a conceptual organising framework built from the reviewed literature, not a validated engineering architecture, and no implementation or evaluation of the framework as a whole is claimed or has been attempted in this work.

Such an architecture would need, at minimum, three coordinated layers. A waveform layer (FHSS, DSSS, or a hybrid) provides baseline communication resilience at the millisecond, single-symbol-transmission timescale. A channel-security layer (PLS via artificial noise, beamforming, or cooperative jamming) operates at the packet or frame timescale, reshaping the channel so as to reduce the eavesdropper's received signal quality while preserving the legitimate link. An adaptive-intelligence layer, operating at the seconds-to-minutes timescale of policy execution and learning, monitors the electromagnetic environment, recognises threats, and adjusts the parameters passed down to the two layers below.

The key architectural requirement is bidirectional communication between these layers. The adaptive-intelligence layer needs accurate, timely feedback from the waveform layer on link quality and interference. The channel-security layer needs position and trajectory information from the swarm's coordination system to optimise beamforming and jamming geometry. The waveform layer needs updated hopping sequences and spreading parameters from the cognitive layer faster than the jammer's own reaction time. All three layers draw on the same limited power, bandwidth, and compute budget, and none of the frameworks reviewed here explicitly defines the interfaces between layers, what information each layer needs from the others, or how conflicting resource demands between layers should be arbitrated. Table 6 sets out the open challenges implied by this conceptual gap, organised by layer and priority.

Table 5. Taxonomy of security layers, techniques, and threats addressed

Layer	Technique	Primary Threat Addressed
Waveform Layer	FHSS - Frequency Hopping Spread Spectrum	Active jamming (narrowband, reactive)
Waveform Layer	DSSS - Direct Sequence Spread Spectrum	Active jamming (broadband barrage)
Waveform Layer	Hybrid FHSS/DSSS	Active jamming (reactive + broadband combined)
Waveform Layer	QRNG-based FHSS	Predictive/pattern-learning reactive jamming
Channel Security Layer	Artificial noise injection	Passive eavesdropping
Channel Security Layer	Cooperative friendly jamming	Passive eavesdropping in swarm architectures
Channel Security Layer	Beamforming / spatial nulling	Eavesdropping + directional jamming
Channel Security Layer	Intelligent Reflecting Surfaces (IRS)	Eavesdropping (channel reshaping)
Adaptive Intelligence Layer	DRL-based cognitive radio	Adaptive/intelligent jamming
Adaptive Intelligence Layer	MARL with CTDE	Coordinated swarm-level adaptive jamming
Adaptive Intelligence Layer	Game-theoretic / Stackelberg	Strategic adversarial interaction
Adaptive Intelligence Layer	DRL + energy harvesting	Adaptive jamming + energy sustainability

6.5 Open challenges and a future research roadmap

Table 6 organises the open challenges identified in this review according to the layer they mainly affect and by research priority. This prioritisation reflects which challenges most directly affect military deployability, based on the analysis in Sections 2–5. It is not intended as a judgement on how intrinsically difficult a given challenge is.

A related and very recent line of work shows one plausible route toward closing Gap 4 (cross-layer optimisation) at the physical layer specifically. Silva and Grilo [38] jointly optimise UAV swarm formation, beam-steering antenna direction, and traffic routing using a genetic algorithm. They report that coordinating swarm-level formation behaviour with antenna-level beam steering clearly outperforms either mechanism applied on its own,

though at a computational cost the authors themselves identify as a barrier to real-time use. The approach is preliminary and centralised rather than distributed, and the authors point to decentralised, MARL-based extensions as future work. Even so, it is a concrete, 2026 example of the kind of joint waveform-and-swarm-behaviour optimisation this review argues is largely missing from the literature (Section 6.2, Gap 4).

Table 6. Open challenges and future research roadmap

#	Challenge	Description	Layer	Priority
1	Integration architecture	No validated framework specifies how waveform, PLS, and cognitive layers exchange information and share resources jointly	All three	Critical
2	Adversarial co-adaptation	DRL defenders are evaluated against static jammers; co-trained adversarial evaluation is absent from the literature	Cognitive radio	Critical
3	Operational-scale MARL	Published MARL results involve ≤12 UAVs; scalability to 50–200-node swarms is unvalidated	Cognitive radio	High
4	LoS PLS under mobility	PLS theory assumes a channel-quality advantage; UAV LoS geometry systematically degrades this in open terrain	PLS	High
5	Sensing reliability at low SNR	Jammer classification accuracy degrades sharply below 5 dB SNR; real environments are frequently noisier than benchmark conditions	Cognitive radio	High
6	Key management at swarm scale	Shared-key FHSS/DSSS face distribution, rotation, and revocation challenges across large autonomous swarms	Waveform	Moderate
7	Resource arbitration across layers	Artificial noise, beamforming, and spectrum sensing compete for the same power/bandwidth/compute pool	PLS + Cognitive	Moderate
8	Graceful degradation	No framework specifies how the swarm behaves when one security layer is selectively jammed or overloaded	All three	Moderate

Beyond these near-term challenges, generative-AI-based security frameworks represent a longer-horizon research direction. Zolfaghari et al. [32] discuss how generative AI could support adaptive communication semantics for UAV swarms operating in space-air-ground integrated networks, allowing swarms to adapt semantics as well as waveforms under adversarial pressure. Alsadie [33] separately discusses adversarial AI training, self-healing AI models, and blockchain-secured network management as complementary parts of defence-in-depth against the full range of cyber threats a military UAV swarm may face. Both directions remain largely conceptual at this stage, but they suggest that the layered-architecture framework proposed in Section 6.4 will eventually need to accommodate AI-native security primitives that are still at an early research stage.

6.5. Technology readiness: practicality against deployment

A systematic review of UAV swarm anti-jamming techniques is not complete without also assessing how close each approach actually is to deployment. Because academic literature naturally favours algorithmic novelty, it can create an impression of field readiness that is more optimistic than a systems-engineering perspective would support. Table 7 gives a technology-readiness assessment for the major technique families reviewed here, citing SDR or hardware results where they exist.

The pattern shown in Table 7 has a direct engineering implication: in this literature, algorithmic complexity and deployment readiness move in opposite directions. FHSS and DSSS are fully deployed; PLS and DRL remain largely experimental; MARL has not moved beyond simulation. This is not really surprising, given how research incentives work - researchers are rewarded for pushing performance ceilings, not for closing deployment gaps - but it leaves the defence engineering community with a real design choice: build around proven but limited waveform-level techniques, or accept the risk of fielding less mature but more capable AI-driven techniques. A

third path, deploying mature techniques now while building a disciplined, repeatable validation pipeline for the more advanced techniques, has received relatively little attention in the academic literature reviewed here, but it is arguably the more defensible engineering position.

Three structural reasons explain why AI-based techniques remain mostly simulation-based.

First, DRL and MARL policies typically need thousands of training episodes to converge, and UAV platforms cannot currently perform this training onboard. Policies are trained offline and then uploaded, which in turn requires a standardised, sufficiently realistic training environment that does not yet exist across the diversity of real electromagnetic conditions.

Second, running a trained DRL policy in real time exceeds the onboard compute available on many of the small platforms most likely to make up large swarms, and inference cost scales with the number of coordinating agents.

Third, SDR platforms capable of evaluating DRL policies in real radio environments at operationally relevant update rates are expensive and specialised, which limits simulation-to-hardware validation to a small number of research groups with the resources to bridge this gap. These are engineering and resourcing challenges rather than fundamental theoretical barriers, but they explain the size of the technology-readiness gap and the scale of investment needed to close it.

Table 7. Technology readiness of anti-jamming techniques for UAV swarm communications

Technique	Simulation	SDR	Hardware Prototype	Field Test	Operational Status
FHSS	✓	✓	✓	✓	Mature (deployed)
DSSS	✓	✓	✓	✓	Mature (deployed)
Hybrid FHSS/DSSS	✓	Limited	Limited	Rare	Developing
PLS -Artificial Noise	✓	Rare	No	No	Experimental
PLS - IRS	✓	No	No	No	Emerging
DRL Cognitive Radio	✓	Few	Few	No	Experimental
MARL	✓	No	No	No	Early Research

6.6. Research gaps and engineering challenges

The analysis in Sections 2–5 identifies eight engineering gaps between the current research literature and a deployable, operationally resilient communication-security architecture for military UAV swarms. These are presented according to the layer they mainly affect, although several of them are cross-layer in practice.

Gap 1: Insufficient Realistic Channel Models. Most reviewed studies rely on simplified propagation assumptions - idealised Rician or Rayleigh fading, stationary UAV positions, and uniform scattering - that do not capture the rapidly changing, geometry-dependent channel conditions a mobile swarm actually experiences. Performance metrics derived under these assumptions are hard to validate operationally, and no standardised, high-fidelity channel model exists against which all three security layers could be consistently evaluated. Closing this gap will require coordinated investment between the channel-modelling and anti-jamming research communities.

Gap 2: Limited Experimental Validation. Table 7 makes this explicit: the strongest current anti-jamming techniques - DRL, MARL, and PLS with artificial noise - have been validated almost exclusively in simulation, for narrow UAV configurations. Only a small number of research groups report SDR-based validation at the scale and update rate that cognitive anti-jamming needs, and hardware-in-the-loop or field testing is essentially absent from the reviewed literature. This simulation-to-SDR transition should be treated as the single most

critical stage in the technology-readiness pipeline, and journal and conference programmes could do more to encourage the experimental work needed to bridge it.

Gap 3: High Onboard Computational Cost of AI Techniques. DRL and MARL bring inference and coordination overheads that many current UAV platforms, particularly the small, lightweight platforms typical of large swarms, cannot support onboard. This is a deployment-limiting factor, not merely an implementation detail, and the existing literature leans toward software-level algorithm design at the expense of the hardware side of the problem - embedded GPUs, FPGAs, and neuromorphic processors among them.

Gap 4: Absence of Cross-Layer Optimisation. Most reviewed papers optimise a single mechanism - a waveform parameter, a PLS policy, or a cognitive-radio strategy - without accounting for the other mechanisms it must coexist with in a deployed system. The resource competition described in Section 6.2 cannot be resolved within a single-layer optimisation framework. A genuinely integrated architecture needs cross-layer resource allocation treated as one single optimisation problem, which is largely missing from the current literature.

Gap 5: Lack of Energy-Aware Anti-Jamming Design. Every security measure reviewed here draws on a UAV's limited power budget, shared with propulsion, sensors, and payload. Very few studies in this review optimise the energy consumption of the full security stack jointly with communication-security performance, yet doing so is a prerequisite for realistic deployment, not an optional refinement.

Gap 6: Unvalidated MARL Scalability. Published MARL results are almost always demonstrated on swarms of four to twelve agents. Moving to fifty- to two-hundred-node swarms substantially increases intra-swarm interference, coordination overhead, and the difficulty of maintaining distributed consistency and training convergence - exactly the regime military swarms are likely to operate in. Methods for validating MARL scaling into this regime do not yet exist, and the underlying theory of convergence and robustness for large, dynamic multi-agent systems under adversarial conditions remains immature. This may be the single most consequential open problem among the AI-related gaps identified in this review.

Gap 7: Absence of Standardised Benchmarking. The literature reviewed here uses a wide range of simulation environments, channel models, jammer models, and performance metrics, which makes direct cross-study comparison difficult and, in some places, unreliable. A common benchmark - standardised channel models, jammer models, swarm configurations, and metrics - comparable to benchmarking efforts in other applied machine learning fields, would substantially improve confidence in cross-study conclusions and help direct future hardware-validation effort toward the most promising approaches.

Gap 8: No Framework for Graceful Degradation. Military swarm systems need to remain operational when one or more security layers is degraded, overloaded, or selectively jammed. None of the reviewed literature specifies how a layered architecture should behave if the cognitive-radio layer's accuracy falls below what is needed, if too few platforms remain for PLS coverage, or if individual platforms fail. Graceful degradation - reducing security posture without losing essential capability - is a basic requirement for any system meant to operate in a contested environment, and it is not addressed in the literature reviewed here. It should be treated as a first-order requirement in future system-level design work.

7. Conclusions

This systematic review examined 35 studies covering three related dimensions of secure UAV swarm communication: waveform-level defence using FHSS and DSSS, interception resistance based on Physical Layer Security, and adaptive countermeasures through AI-based cognitive radio and multi-agent reinforcement learning. None of these three research directions, alone or combined, can be considered a solved engineering problem. Each one shows real and well-documented progress, but none has managed to close the gap between simulated performance and what a military swarm could actually depend on in the field against an adaptive, AI-capable adversary. This gap, rather than the merit of any single technique, is the main finding of the present review.

FHSS and DSSS remain the operational baseline mainly because they are simple, mature, and well understood, and this also means that their failure modes are relatively predictable. Open-source reporting from the conflict

in Ukraine, considered together with the limitations discussed in Section 3.3, agrees with the literature in showing that predictable hopping patterns perform poorly against reactive jammers [5, 7]. The processing-gain trade-off of DSSS is also difficult to reconcile with the latency requirements of multi-UAV coordination [8]. Both techniques additionally share a structural dependency on a shared cryptographic key, which needs to be distributed and protected across the whole swarm [10]. This does not reduce the practical usefulness of spread-spectrum techniques; it only defines the conditions under which they should be evaluated before additional capability is built on top of them.

For PLS, the main obstacle is not the information-theoretic security concept itself, which does not require key distribution as spread-spectrum techniques do, but the line-of-sight-dominated channel geometry typical of low-altitude UAV operation. This condition violates one of the working assumptions behind most PLS analysis reported so far [16, 17, 31]. As a result, trajectory optimisation and artificial-noise injection become more difficult to achieve for UAV swarms than for the terrestrial scenarios for which PLS theory was originally developed. Both approaches also consume energy and computing resources that are limited within a swarm, and the coordination overhead at large, mobile swarm scale is still largely unaddressed.

For AI-driven cognitive radio and MARL, the simulation results reported in the literature, 5–12 dB SINR improvement and 30–60% throughput improvement over static baselines, are large enough to be operationally significant against an adversary that now regularly uses adaptive electronic warfare. Whether these gains can actually be achieved in real operation depends on the three gaps identified in Section 5.5: unreliable sensing under realistic SNR conditions, the vulnerability of the learning process to adversarial manipulation, and the absence of validated results at swarm scale. The evaluation protocol commonly used in this literature, which tests a learned defender against a fixed jammer benchmark, does not really answer any of these three questions.

The overall argument of this review is that none of these three techniques represents, by itself, the decisive open problem in UAV swarm communication security. The real open problem is architectural in nature: the three layers depend on one another, must exchange information across different timescales, need to negotiate shared resource constraints, and the overall system should degrade gracefully instead of failing completely when one layer is compromised.

The layered framework introduced in Section 6.4 is intended as a conceptual model of this interdependence, a way of organising the problem, rather than as a validated system design. Section 6.6 showed that technology readiness decreases as algorithmic sophistication increases, with MARL still limited to simulation and DRL only reaching early SDR validation. Section 6.7 identified eight specific engineering gaps, ranging from channel-model realism and experimental validation to cross-layer optimisation, standardised benchmarking, and graceful degradation, that a truly deployable architecture would still need to address.

Three recommendations can be made for the defence engineering community responsible for specifying, procuring, and deploying UAV swarm communication systems. First, candidate systems should be evaluated at the level of the complete security architecture rather than at the level of individual components, since component-level performance figures do not combine in a predictable way (Section 6.2).

Second, adversarial robustness should be treated as a first-order design requirement rather than a post-hoc validation step: co-adaptive adversarial training should be a standard part of developing any anti-jamming AI system intended for military use, not an optional addition. Third, research and engineering effort should be directed toward the interfaces between layers, that is, how the waveform, channel-security, and cognitive-adaptation layers negotiate information and resources, since this is the area least addressed by the current literature, and arguably the one where progress would matter most.

The research community is actively working on the open questions surveyed in this review, and the individual techniques discussed here represent genuine, well-supported progress. What is still missing is an integrated, adversarially validated, swarm-scale architecture that a military operator could deploy with confidence. Closing this gap will require sustained cooperation between the research and defence-engineering communities, and the priorities outlined above are offered as a starting point for that effort.

Declaration of competing interest

The author declares that he has no known financial or non-financial competing interests in any material discussed in this paper.

Funding information

No funding was received from any financial organisation to conduct this research.

Author contribution

M. H. Khalil: conceptualisation, methodology, literature search and screening, quality assessment, analysis, and writing - original draft, review, and editing. The author approved the final version of the manuscript.

Ethical approval statement

Ethical approval is not applicable for this research, as this is a systematic literature review that did not involve human or animal subjects.

Informed consent

Informed consent is not applicable for this research, as this manuscript does not contain personal data, images, or other information requiring participant consent.

Declaration of use of AI in the writing process

A generative AI tool (Claude, Anthropic) was used only for formatting the manuscript to the journal template. No AI-generated content was included in the scientific text. The author reviewed and approved the final manuscript.

References

[1] J. Tao, Y. Shi, Y. Li, L. Wan, and Y. Zhang, "Multi-slot optimized index modulation based on FHSS for efficient anti-jamming," Physical Communication, vol. 71, Art. no. 102657, Aug. 2025, doi: 10.1016/j.phycom.2025.102657.

[2] Kang, J.G., and Choi, B.C. "Multi-modem-based FHSS-drone takeover with precision spoofing." ETRI Journal 47, no. 3 (2025): 410–421. https://doi.org/10.4218/etrij.2024-0369

[3] Z. Spasojevic and J. Burns, "Performance comparison of frequency hopping and direct sequence spread spectrum systems in the 2.4 GHz range," in 11th IEEE International Symposium on Personal, Indoor and Mobile Radio Communications (PIMRC), London, UK, 2000, pp. 426–430, doi: 10.1109/PIMRC.2000.881461.

[4] Peng, J., Zhang, Z., Wu, Q., and Zhang, B. "Anti-Jamming Communications in UAV Swarms: A Reinforcement Learning Approach." IEEE Access 7 (2019): 180532–180543, doi: 10.1109/ACCESS.2019.2958328.

[5] M. Adil, M. A. Jan, Y. Liu, H. Abulkasim, A. Farouk, and H. Song, "A Systematic Survey: Security Threats to UAV-Aided IoT Applications, Taxonomy, Current Challenges and Requirements with Future Research Directions," IEEE Transactions on Intelligent Transportation Systems, vol. 24, no. 2, pp. 1437–1455, Feb. 2023, doi: 10.1109/TITS.2022.3220043.

[6] J. de Curtò, I. de Zarzà, J.-C. Cano, and C. T. Calafate, "Enhancing Communication Security in Drones Using QRNG in Frequency Hopping Spread Spectrum," Future Internet, vol. 16, no. 11, p. 412, Nov. 2024, doi: 10.3390/fi16110412.

[7] Lv, Z., Xiao, L., Du, Y., Niu, G., Xing, C., and Xu, W. "Multi-Agent Reinforcement Learning Based UAV Swarm Communications Against Jamming." IEEE Transactions on Wireless Communications 22, no. 12 (2023): 9063–9075. https://doi.org/10.1109/TWC.2023.3268082

[8] B. Wang, J. Fang, J. Du, and S. Shao, "Jamming-Resistant UAV Communications: A Multichannel-Aided Approach," IEEE Wireless Communications Letters, to be published. [Preprint]. Available: arXiv:2507.14945.

[9] Van Huynh, N., Hoang, D.T., Nguyen, D.N., and Dutkiewicz, E. "DeepFake: Deep Dueling-based Deception Strategy to Defeat Reactive Jammers." IEEE Transactions on Wireless Communications 20, no. 10 (2021): 6898–6914. https://doi.org/10.1109/TWC.2021.3078439

[10] Wei, X., Ma, J., and Sun, C. "A Survey on Security of Unmanned Aerial Vehicle Systems: Attacks and Countermeasures." IEEE Internet of Things Journal 11, no. 21 (2024): 34826–34847. https://doi.org/10.1109/JIOT.2024.3429111

[11] B. Ji, Y. Liu, T. Li, L. Xing, W. Wang, S. Mumtaz, X. Shang, W. Liu, and C. Han, "Anti-Jamming Null Space Projection Beamforming Based on Symbiotic Radio," CMES — Computer Modeling in Engineering and Sciences, vol. 138, no. 1, pp. 679–689, 2024, doi: 10.32604/cmes.2023.028667.

[12] C. O. Nnamani, M. R. A. Khandaker, and M. Sellathurai, "UAV-Aided Jamming for Secure Ground Communication with Unknown Eavesdropper Location," IEEE Access, vol. 8, pp. 72881–72892, 2020, doi: 10.1109/ACCESS.2020.2986025.

[13] H. Dang-Ngoc, D. N. Nguyen, K. Ho-Van, D. T. Hoang, E. Dutkiewicz, Q.-V. Pham, and W.-Joo Hwang, "Secure Swarm UAV-Assisted Communications with Cooperative Friendly Jamming," IEEE Internet of Things Journal, vol. 9, no. 24, pp. 25596–25611, Dec. 2022, doi: 10.1109/JIOT.2022.3197975.

[14] J. Wang, X. Wang, R. Gao, C. Lei, W. Feng, N. Ge, S. Jin, and T. Q. S. Quek, "Physical Layer Security for UAV Communications in 5G and Beyond Networks," 2021, arXiv:2105.11332. (Remains unpublished in a peer-reviewed venue as of this review; cited by subsequent peer-reviewed work, e.g. J. Aerospace Information Systems, 2024.)

[15] X. Wang, W. Feng, Y. Chen, and N. Ge, "UAV Swarm-Enabled Aerial CoMP: A Physical Layer Security Perspective," IEEE Access, vol. 7, pp. 120901–120916, 2019, doi: 10.1109/ACCESS.2019.2936680.

[16] G. Zhang, Q. Wu, M. Cui, and R. Zhang, "Securing UAV Communications via Trajectory Optimization," in GLOBECOM 2017 - 2017 IEEE Global Communications Conference, Singapore, 2017, pp. 1–6, doi: 10.1109/GLOCOM.2017.8254971.

[17] G. Zhang, Q. Wu, M. Cui, and R. Zhang, "Securing UAV Communications via Joint Trajectory and Power Control," IEEE Transactions on Wireless Communications, vol. 18, no. 2, pp. 1376–1389, Feb. 2019, doi: 10.1109/TWC.2019.2892461.

[18] A. Krayani, A. S. Alam, L. Marcenaro, A. Nallanathan, and C. Regazzoni, "An Emergent Self-Awareness Module for Physical Layer Security in Cognitive UAV Radios," IEEE Transactions on Cognitive Communications and Networking, vol. 8, no. 2, pp. 888–906, Jun. 2022, doi: 10.1109/TCCN.2022.3161937.

[19] H. Yang, Y. Liu, X. Li, Z. Bai, L. Yang, G. Pan, and H. Liu, "Physical layer security and covert communication in UAV-ISAC networks: A comprehensive survey," Journal of King Saud University - Computer and Information Sciences, vol. 37, no. 10, Art. no. 312, 2025, doi: 10.1007/s44443-025-00291-0.

[20] X. Li, Y. Xiang, J. Zhou, Y. Luo, Q. Du, D. Hou, and J. Tian, "Physical-Layer Security Enhancement for UAV Downlink Communication Using Joint Precoding and Artificial Noise Design in Generalized Spatial Directional Modulation," Drones, vol. 9, no. 3, p. 199, Mar. 2025, doi: 10.3390/drones9030199.

[21] W. U. Khan, E. Lagunas, Z. Ali, M. A. Javed, M. Ahmed, and S. Chatzinotas, "Opportunities for Physical Layer Security in UAV Communication Enhanced with Intelligent Reflective Surfaces," IEEE Wireless Communications, vol. 29, no. 6, pp. 22–28, Dec. 2022, doi: 10.1109/MWC.001.2200125.

[22] J. Xue, J. Xu, X. Guan, H. Zhang, and Y. Gan, "Secure and energy-efficient transmission in UAV-assisted intelligent reflecting surface networks," Scientific Reports, vol. 15, no. 1, Art. no. 34538, Oct. 2025, doi: 10.1038/s41598-025-17852-y.

[23] X. Feng, H. Wen, R. Zhao, T. Tang, W. Shi, and Y. Peng, "Communication Anti-jamming System Based on Deep Reinforcement Learning," in Proceedings of the 2023 International Conference on Wireless Communications, Networking and Applications (WCNA 2023), (Lecture Notes in Electrical Engineering, vol. 1361), P. Siarry, M. A. Jabbar, S. K. S. Cheung, and X. Li, Eds. Singapore: Springer, 2025, pp. 126–133, doi: 10.1007/978-981-96-2409-6_13.

[24] A. Krayani, A. S. Alam, L. Marcenaro, A. Nallanathan, and C. Regazzoni, "A Novel Resource Allocation for Anti-Jamming in Cognitive-UAVs: An Active Inference Approach," IEEE Communications Letters, vol. 26, no. 10, pp. 2272–2276, Oct. 2022, doi: 10.1109/LCOMM.2022.3190971.

[25] Z. Yin, J. Li, Z. Wang, Y. Qian, Y. Lin, F. Shu, and W. Chen, "UAV Communication Against Intelligent Jamming: A Stackelberg Game Approach with Federated Reinforcement Learning," IEEE Transactions on Green Communications and Networking, vol. 8, no. 4, pp. 1796–1808, Sep. 2024, doi: 10.1109/TGCN.2024.3373886.

[26] L. Xiao, Y. Ding, J. Huang, S. Liu, Y. Tang and H. Dai, "UAV Anti-Jamming Video Transmissions With QoE Guarantee: A Reinforcement Learning-Based Approach," IEEE Transactions on Communications, vol. 69, no. 9, pp. 5933–5947, Sept. 2021, doi: 10.1109/TCOMM.2021.3087787.

[27] Imran, M., Ibrahim, K., Zhiwen, P. et al. "Particle Swarm Optimized Deep Learning for Jamming Detection and Throughput Enhancement in Cognitive Radio Networks." Scientific Reports (2026). https://doi.org/10.1038/s41598-026-41642-9

[28] S. Yang, "Analysis Of Deep Learning-Based Anti-Jamming Method for UAV Communication," Highlights in Science, Engineering and Technology, vol. 103, pp. 246–252, 2024, doi: 10.54097/kkfwre06.

[29] R. Dong, B. Wang, J. Tian, T. Cheng and D. Diao, "Deep Reinforcement Learning Based UAV for Securing mmWave Communications," IEEE Transactions on Vehicular Technology, vol. 72, no. 4, pp. 5429–5434, April 2023, doi: 10.1109/TVT.2022.3224959.

[30] Wang, X., et al. "A Survey on Security of UAV Swarm Networks: Attacks and Countermeasures." ACM Computing Surveys 57, no. 3 (2024): 1–37. https://doi.org/10.1145/3703625

[31] Y. Zhou et al., "Improving Physical Layer Security via a UAV Friendly Jammer for Unknown Eavesdropper Location," IEEE Transactions on Vehicular Technology, vol. 67, no. 11, pp. 11280–11284, Nov. 2018, doi: 10.1109/TVT.2018.2868944.

[32] Zolfaghari, B., Abbasmollaei, M., Hajizadeh, F., Yanai, N., and Bibak, K. "Secure UAV (Drone) and the Great Promise of AI." ACM Computing Surveys 56, no. 11 (2024): Article 282. https://doi.org/10.1145/3673225

[33] Alsadie, D. "Cybersecurity and Artificial Intelligence in Unmanned Aerial Vehicles: Emerging Challenges and Advanced Countermeasures." IET Information Security (2025). https://doi.org/10.1049/ise2/2046868

[34] Nawaz, T., and Alzahrani, A. "Machine-Learning-Assisted Cyclostationary Spectral Analysis for Joint Signal Classification and Jammer Detection at the Physical Layer of Cognitive Radio." Sensors 23, no. 16 (2023): 7144. https://doi.org/10.3390/s23167144

[35] Zeng, L., Wang, W., and Zuo, W. "A Federated Learning Latency Minimisation Method for UAV Swarms Aided by Communication Compression and Energy Allocation." Sensors 23, no. 13 (2023): 5787. https://doi.org/10.3390/s23135787

[36] Page, M. J., McKenzie, J. E., Bossuyt, P. M., Boutron, I., Hoffmann, T. C., Mulrow, C. D., et al. "The PRISMA 2020 statement: an updated guideline for reporting systematic reviews." BMJ 372 (2021): n71. https://doi.org/10.1136/bmj.n71

[37] J. Yang, M. Cui, H. Zhang, F. Ji, Z. Lai, and Y. Wang, "Agent-Based Anti-Jamming Techniques for UAV Communications in Adversarial Environments: A Comprehensive Survey," arXiv:2508.11687, 2025.

[38] T. Silva and A. Grilo, "Anti-Jamming based on Beam-Steering Antennas and Intelligent UAV Swarm Behavior," arXiv:2510.07292, 2026.