MANOHAR PARRIKAR
idsa
MANOHAR PARRIKAR INSTITUTE FOR DEFENCE STUDIES AND ANALYSES
मनोहर पर्रिकर रक्षा अध्ययन एवं विश्लेषण संस्थान

MP-IDSA Issue Brief

Autonomous Electromagnetic Dominance at Sea

Vishnu B. Nair

November 12, 2026

Summary

The contest for the EM spectrum has moved from the physics of jamming to the speed of decision. An edge-deployed, multi-agent architecture that senses with a digital array, reconstructs with compressive sensing and classifies with deep learning, decides with a domain language model, and attacks with reinforcement learning can close the gap between detecting an emitter and countering it in crowded littoral waters.

Introduction

Control of the electromagnetic (EM) spectrum is no longer a supporting function of naval operations. It is the condition on which they rest. The same spectrum carries the sensors, communications and networks on which action in the sea, land, air, space and cyber domains depends, and a force that loses it fights blind and mute. Yet the maritime environments in which navies now expect to operate, such as the approaches to choke points, disputed littoral waters, and the crowded sea lanes of the Indian Ocean Region, are dense with emitters, many of them civilian and many reconfigurable in software. The volume and agility of the Radio Frequency (RF) traffic they generate have outrun both the rule-based Electronic Warfare (EW) suites built to interpret it and the human operators who supervise them.[1]

This brief argues that the gap is integrative rather than inventive. Classification, coordination, distributed attack and model robustness have each been demonstrated, and each is represented in the literature. What does not exist is an architecture that holds them together under the conditions the maritime environment imposes: intermittent connectivity, severe size, weight, power and cost (SWaP-C) limits, and an adversary who may attack the models themselves. What follows describes that architecture, the distributed platforms that carry it, the safeguards that keep it trustworthy and what India should do about it.

The Maritime Electromagnetic Battlespace

Spectral congestion in the littoral is structural, not accidental. The International Telecommunication Union (ITU) allocates the spectrum among services through its Table of Frequency Allocations, and national allocations, including India's National Frequency Allocation Plan (NFAP), are drawn from it. Interference nonetheless arises at the edges of that framework. Many bands are shared on a primary or secondary basis, transmitters radiate beyond their assigned bandwidth through sidebands and out-of-band emissions, and a strong signal in an adjacent band can saturate a receiver tuned elsewhere. Military emissions add further complication, since not every military system operates within a civil allocation.[2]

The margin is narrower than it looks. In the United States, 5G transmissions in the 3.7–3.98 GHz band were found to interfere with aircraft radar altimeters operating at 4.2–4.4 GHz, a separation of only about 220 MHz. India's C-band allocation of

[1] Alida Monica Doriana Barbu, **"The Cognitive Electronic Warfare in the Age of Artificial Intelligence"**, *Strategies XXI*, no. 15, 2024; Nicolae Sfetcu, *Electronic Warfare and Artificial Intelligence*, MultiMedia Publishing, Craiova, 2024.

[2] Stéphane Ricciardi and Cédric Souque, **"Modern Electromagnetic Spectrum Battlefield"**, *PRISM*, Vol. 9, No. 3, 2021.

3.3–3.67 GHz leaves a gap of about 530 MHz, and no comparable incidents have been reported here. Even so, the Directorate General of Civil Aviation (DGCA) has issued an airworthiness directive requiring radio-altimeter tolerance in Indian airspace, which treats the margin as a safety matter rather than dismissing it.[3] A warship's electronic support receiver, asked to work in the same bands under wartime conditions, has no such protection.

The emitters themselves have changed. Hardware-defined, single-function radars have given way to software-defined Active Electronically Scanned Array (AESA) systems that can alter their operating parameters pulse by pulse, and can search, track, guide a missile and jam at the same time.[4] They radiate Low Probability of Intercept (LPI) and Low Probability of Detection (LPD) waveforms such as frequency-agile carriers, non-linear frequency modulation, pseudo-random pulse repetition interval stagger and ultra-wideband spread spectrum, distributing their energy across a wide bandwidth or below the noise floor so that a conventional receiver never registers them at all.

Against this, legacy electronic support measures are federated and hardware-centric. They capture RF energy, down-convert it, extract pulse width, time of arrival, frequency and angle of arrival, and match the result against a static, pre-programmed Electronic Order of Battle (EOB).[5] When a waveform is not in the library, two failures follow. The suite ignores the threat, or it forces the signal into the nearest entry and treats it as that emitter. The second is the more dangerous. A missile-guidance radar catalogued as a surveillance radar draws a countermeasure that has no effect, while the operator sees a track that appears to have been dealt with. Because threat libraries are built by offline laboratory analysis and updated by hand, no such system can respond to a zero-day or dynamically reconfigured waveform in real time.

A platform may face dozens of simultaneous emitters from surface vessels, aircraft and incoming anti-ship missiles, and the Pulse Descriptor Words (PDW) generated by modern wideband receivers saturate both the processing buffers and the operator.[6] Legacy suites do offer automatic modes for countermeasure selection, power assignment and effectiveness assessment, but those modes execute pre-programmed responses drawn from the same static library. With anti-ship missiles

---

[3] ET Infra Desk, **"DGCA Sets 2027 Deadline for Aircraft Upgrades to Address 5G Interference"**, *ETInfra*, 9 April 2026.

[4] David L. Adamy, *EW 101: A First Course in Electronic Warfare*, Artech House, Norwood, MA, 2001; Purabi Sharma, Kandarpa Kumar Sarma and Nikos E. Mastorakis, **"Artificial Intelligence Aided Electronic Warfare Systems — Recent Trends and Evolving Applications"**, *IEEE Access*, Vol. 8, 2020.

[5] Karamtot Krishna Naik, **"Simulation of Cognitive Electronic Warfare System With Sine and Square Waves"**, *Defence Science Journal*, Vol. 73, No. 4, 2023.

[6] Nicolae Sfetcu, *Electronic Warfare and Artificial Intelligence*, no. 1.

at supersonic or hypersonic speed, the response window falls to a fraction of a second, which is too short for a human to intervene and too short to correct an automatic mode that has chosen wrongly.

The operating environment compounds the problem. A contested EM environment implies intermittent or total loss of satellite and shore connectivity, whether from jamming or from cyber attack.[7] Forward units must therefore work on the principle of mission command, with authority delegated downward. The EW suite must become an island of intelligence at the edge, perceiving, deciding, and acting on local computation alone.

Table 1. The Cognitive Gap in Maritime EW

Existing approach	What it delivers	What remains unresolved
Rule-based EW with static threat libraries	Fast, predictable and verifiable responses to emitters already catalogued	Cannot classify or counter a waveform absent from the library; cannot adapt without a manual update
Machine-learning signal classification	High accuracy in modulation classification and emitter identification, including at low signal-to-noise ratio	Treats classification as an end in itself; no closed loop from classification to countermeasure selection
Multi-agent orchestration with large language models	Coordination of distributed EW assets and reasoning over tactical options	Assumes reliable connectivity and substantial computing power; untested against adversarial manipulation
UAV swarm electronic warfare	Distributed jamming power; relief from the SWaP-C limits of individual platforms	Addresses the attack function alone; no integrated sense–decide–act loop
Adversarial-AI defence and explainability	Model robustness and transparency, considered in isolation	Not embedded in an operational EW architecture; no defined handover of authority to a human operator

Source: Author's synthesis of the literature surveyed in this brief.

[7] Stéphane Ricciardi and Cédric Souque, **"Modern Electromagnetic Spectrum Battlefield"**, no. 2, p. 136.

No single row of the table is the problem. Classification without a closed loop to the countermeasure, orchestration that assumes connectivity it will not have, jamming without sensing and robustness designed in isolation — each solves one part of the problem while assuming the rest are solved elsewhere. The gap lies between the rows. Each gap has a corresponding answer in the design. The ES Agent's compressive-sensing receiver and its convolutional and recurrent classifiers handle waveforms absent from the library, working from the signal itself rather than a stored template.

The missing link from classification to countermeasure is closed by the Decision and EA Agents, which turn a classification into a selected technique and then measure its effect. The assumption of reliable connectivity is removed by placing the entire stack at the edge and linking the swarm by mesh. The swarm does not decide on its own: the Decision Agent on the host platform assigns the mission. At the same time, each node executes its assigned role locally and shares track state over the mesh, so that decisions are taken where the latency budget allows and coordinated where it does not. The gap between a robust model and a robust architecture is closed by embedding adversarial training, uncertainty gating and explainability in the same agent structure, with authority returning to the operator wherever confidence falls.

A Multi-Agent Architecture for the Edge

The proposal is a decentralised, edge-deployed multi-agent system integrated into the platform's Combat Management System (CMS).[8] A monolithic model would concentrate failure and add latency; dividing the mission among specialised cooperative agents avoids both. The agents run concurrently on heterogeneous edge hardware and communicate over standard vehicle buses such as the NATO Generic Vehicle Architecture.

The Sensory Agent manages a multi-channel digital antenna array. Where analogue receivers scan slowly and see a narrow instantaneous bandwidth, the array samples RF directly at the element level through high-speed analogue-to-digital converters. The agent runs digital beamforming to form several independent beams at once. When it detects a jammer or missile seeker, it places deep spatial nulls in that direction, shielding its receivers from saturation while keeping sensitive beams elsewhere.

The Electronic Support (ES) Agent turns those samples into identifications. Direct sampling at the antenna produces data at a rate no shipboard processor can handle

[8] Vadym Kozlov, Vadym Slyusar, Volodymyr Tverdokhlibov and Zoia Andriichuk, **"Integration and Coordination of Electronic Warfare Assets Through Large-Scale Language Models"**, *Technology Audit and Production Reserves*, Vol. 1, No. 2(81), 2025.

in full. Since most real signals are sparse, the receiver reconstructs them from far fewer samples than the Nyquist rate would demand, reducing converter power and computational load. Sparse Bayesian Learning performs reconstruction at the edge and does not need to be told in advance which representation makes the signal sparse.[9] The reconstructed signal then passes through a staged pipeline. A one-dimensional convolutional neural network performs automatic modulation classification with over 90 per cent accuracy even below 0 dB Signal-to-Noise Ratio (SNR).[10] Specific Emitter Identification (SEI) maps the signal to a physical platform, and a Long Short-Term Memory (LSTM) network, trained on sequential pulse-interval data, predicts the radar's next frequency hop or pulse repetition mode so that countermeasures can be scheduled in advance.[11]

The Decision Agent is the cognitive core. It runs a domain-specific transformer over a Mixture of Experts (MoE) framework, sized for field conditions rather than for a data centre. A Retrieval-Augmented Generation (RAG) pipeline connected to the platform's local signal library converts an unknown emitter's metadata into vector embeddings and retrieves historical tactics and recommended responses. The mixture routes the threat profile to specialised models: one weighs lethality through fuzzy logic, one schedules countermeasure resources, and one deconflicts distributed assets.[12] The computational weight of this layer is the least settled part of the design. Unlike the ES classifier, whose compression to a fraction of a megabyte is demonstrated in the literature, a Mixture-of-Experts transformer with a vector-search retrieval pipeline and three further sub-models has no comparable published footprint at the edge. Whether it can run within the power, cooling, and space available on a warship is an open question, and a prototype would have to settle it before the architecture could be fielded.

The Electronic Attack (EA) Agent executes. It uses deep reinforcement learning of the policy-gradient kind, in an actor–critic arrangement, where a learned value estimate serves as a baseline that reduces the variance of the policy update.[13] Operating in a continuous action space, it selects waveforms and schedules power and frequency to avoid disrupting friendly communications.[14] It then uses the victim radar's own

[9] Purabi Sharma, Kandarpa Kumar Sarma and Nikos E. Mastorakis, **"Artificial Intelligence Aided Electronic Warfare Systems — Recent Trends and Evolving Applications"**, no. 4; Tanner McWhorter et al., **"Machine Learning Aided Electronic Warfare System"**, *IEEE Access*, Vol. 9, 2021.

[10] Purabi Sharma, Kandarpa Kumar Sarma and Nikos E. Mastorakis, **"Artificial Intelligence Aided Electronic Warfare Systems — Recent Trends and Evolving Applications"**, no. 4.

[11] Tanner McWhorter et al., **"Machine Learning Aided Electronic Warfare System"**, no. 9.

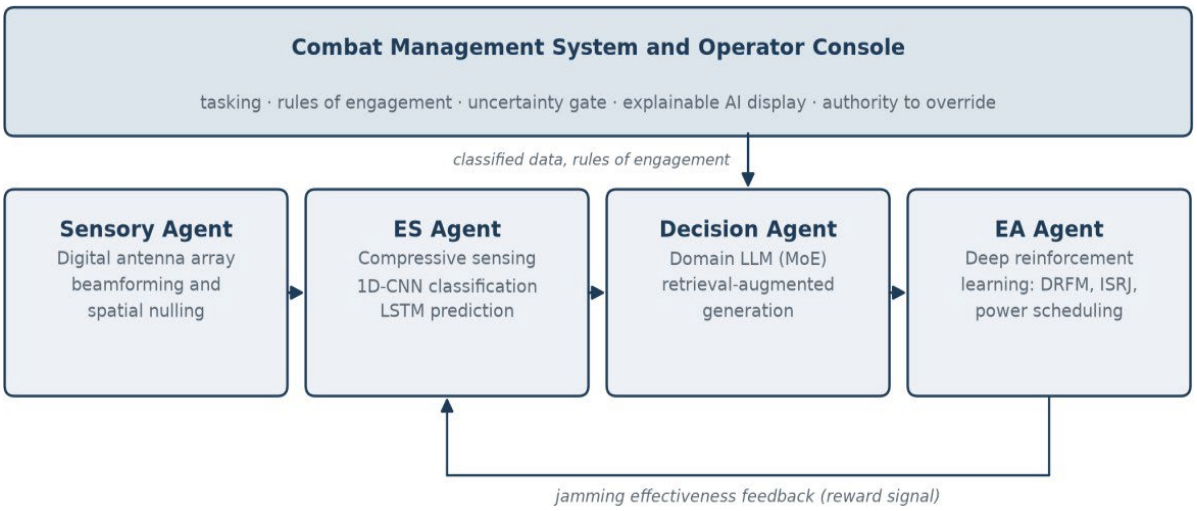
[12] Vadym Kozlov, Vadym Slyusar, Volodymyr Tverdokhlibov and Zoia Andriichuk, **"Integration and Coordination of Electronic Warfare Assets Through Large-Scale Language Models"**, no. 8, pp. 56–57.

[13] Suleiman G. Suleiman, A. E. Airoboman and Nathaniel Dalla, **"Artificial Intelligence-Driven Electronic Warfare Systems: The Future of Spectrum Dominance in Modern Conflict"**, *Journal of Engineering Logic and Modelling_Research*, Vol. 10, No. 5, 2025.

[14] Reeshen Reddy and Saurabh Sinha, **"State-of-the-Art Review: Electronic Warfare Against Radar**

behaviour, losing track or reverting from track to search, as the reward signal that tells it whether the technique worked.[15]

Figure 1. Multi-Agent Architecture of an Edge-Deployed Cognitive EW Suite



Source: Author’s original synthesis, informed by Kozlov and co-authors (2025) and the sources cited throughout this section. The classified element of the suite is not depicted; it sits at the interface between the Decision Agent and the combat management system, where mission data and rules of engagement apply.

Distributed Nodes: UAV Swarms in the Electronic Attack Role

A micro-UAV cannot carry the cooling, generation and aperture needed to jam a multi-megawatt naval air-defence radar. The answer is not a bigger drone but a different division of labour: a distributed autonomous system of dozens of low-cost, expendable nodes, some carrying wideband ES receivers and others low-power software-defined jammers, whose coordinated action synthesises a high-power, multi-directional attack.[16]

Distributed jamming also answers the counter. A single continuous jammer is a beacon for an anti-radiation missile or a passive Home-on-Jam (HoJ) seeker. A swarm can instead ‘flicker in space’, where node A radiates for a few milliseconds and goes silent, node B radiates from a different position, then node C. To the seeker, the emission appears continuous but jumps between physically separated

Systems”, *IEEE Access*, Vol. 13, 2025.

[15] Shafei Wang, Yanfei Bao and Yan Li, **“The Architecture and Technology of Cognitive Electronic Warfare”**, *Scientia Sinica Informationis*, Vol. 48, 2018, pp. 1603–13.

[16] Vadym Slyusar, Vadym Kozlov, Serhii Pochernin and Iryna Nalapko, **“Conceptual Foundations of the Swarm Employment of Unmanned Aerial Vehicles as Intelligent Means of Electronic Warfare”**, *Technology Audit and Production Reserves*, Vol. 3, No. 2(83), 2025.

points, so no stable track forms and the missile passes through the space between the nodes.[17]

None of this survives without a self-repairing network. The swarm forms a local ad hoc mesh and routes with Optimised Link State Routing (OLSR). Raw sensor data moves over the data distribution service, where latency matters most, while command messages and track states go over the lighter Message Queuing Telemetry Transport (MQTT) protocol.[18] The swarm also optimises its geometry, computing how sensitive the jamming effect is to each node's position and moving the node in the direction that has the most effect for the least displacement, so it sits inside the victim radar's main lobe or its more sensitive sidelobes.

Reading the Emitter: Radar Fingerprinting at the Edge

Specific Emitter Identification (SEI) deserves separate treatment, because it changes what electronic support can produce. No two transmitters are identical. Manufacturing tolerances, chemical variance in silicon wafers, microscopic imperfections in metal cavities and differences in transformer windings leave each unit with its own analogue anomalies.[19] The combined effect is Unintentional Modulation on Pulse (UMOP)—microscopic ripples in the envelope and intra-pulse phase that are fixed by the physical geometry of the transmitter, and therefore stable, repeatable and very hard to counterfeit.

That signature does three things a waveform parameter cannot. It exposes decoys and radar spoofing, where an adversary mimics a friendly emission. It holds an accurate EOB even while radars are frequency-agile or operating at low probability of intercept. It converts passive detection into platform-level tracking, letting a warship isolate hostile units from the hundreds of merchant and fishing vessels around it and maintain a targeting-grade passive track.[20]

The classical approach of hand-extracted features such as pulse rise time, settling time and envelope asymmetry is cheap but brittle. Multipath reflection, sea clutter

[17] Reeshen Reddy and Saurabh Sinha, **"State-of-the-Art Review: Electronic Warfare Against Radar Systems"**, no. 14, p. 35562.

[18] Vadym Slyusar, Vadym Kozlov, Serhii Pochernin and Iryna Nalapko, **"Conceptual Foundations of the Swarm Employment of Unmanned Aerial Vehicles as Intelligent Means of Electronic Warfare"**, no. 16, p. 73.

[19] M. A. S. Pathania, **"Machine Learning and RF Fingerprinting for Signal Intelligence and Tactical Communication Systems: A Review"**, *International Journal of Creative and Open Research in Engineering and Management*, Vol. 2, No. 6, 2026.

[20] Maj Ankur Singh Pathania, **"Machine Learning and RF Fingerprinting for Signal Intelligence and Tactical Communication Systems: A Review"**, *International Journal of Creative and Open Research in Engineering and Management*, Vol. 2, No. 6, 2026; Saeif Al-Hazbi, **"Deep-Learning-Based Cross-Layer Fingerprinting for Secure and Trustworthy Communication Systems"**, PhD Thesis, Hamad Bin Khalifa University, 2025.

and atmospheric fading warp the transient envelope, and accuracy falls away precisely at the ranges where the SNR drops below 10 dB. Deep networks such as Convolutional Neural Networks (CNNs) and self-attention transformer encoders fare far better, as they weigh every sample in the sequence against every other.[21] The caveat, however, is data. Accuracy depends on the number of observations and the number of samples per observation, and AI methods still trail classical ones in dense, low-signal conditions. Hybrid approaches are therefore the more promising near-term direction.[22]

More than accuracy, deployment is the binding constraint. Running these models on a warship means processing millions of pulses per second without server racks. Compression under the TinyML framework, such as TensorFlow Lite, closes the gap by quantising weights from 32-bit floating point to 8-bit integers. A baseline convolutional model shrinks from about 1.4 MB to 123.8 KB, and a transformer encoder to 73.27 KB, with no material performance loss.[23]

Securing the Cognitive Suite

A suite that is autonomous but open to manipulation is more dangerous than the rule-based system it replaces, because it will act quickly and confidently on a false picture. Four threat vectors matter. In model evasion, an adversary shapes its own radar pulses so that the classifier returns a confident but wrong label, and the correct self-defence measure is never released. Secondly, in data poisoning, deceptive or mislabeled emissions radiated in peacetime enter the local training set and shift the decision boundary, leaving a trigger waveform that opens a backdoor later. In model inversion, structured probe signals and the countermeasures they draw allow an adversary to reconstruct the network's weights and map its blind spots. Lastly, in supply chain compromise, backdoors enter through open-source libraries or pre-trained foundation models.[24]

These defences applied together are proportionate. The first is RF-specific adversarial training, in which gradient-based perturbations generated by the Fast Gradient Sign Method (FGSM) are mixed into the training set, forcing decision boundaries that generalise against deception. The second is uncertainty-based action gating. Models emit a confidence or entropy measure alongside the classification; when uncertainty

[21] Ahmed Mohamed Hussain, Nada Abughanam and Panos Papadimitratos, **"Edge AI-based Radio Frequency Fingerprinting for IoT Networks"**, arXiv:2412.10553, December 2024.

[22] Karen Zita Haigh and Julia Andrusenko, *Cognitive Electronic Warfare: An Artificial Intelligence Approach*, 2nd ed., Artech House, Norwood, MA, 2025.

[23] Ahmed Mohamed Hussain, Nada Abughanam and Panos Papadimitratos, **"Edge AI-based Radio Frequency Fingerprinting for IoT Networks"**, n. 21.

[24] Mathew J. Walter, Aaron Barrett, David J. Walker and Kimberly Tam, **"Adversarial AI Testcases for Maritime Autonomous Systems"**, *AI, Computer Science and Robotics Technology*, Vol. 2, No. 1, 2023.

crosses a threshold, the suite is barred from autonomous action, falls back to a deterministic, rule-based safe state, and alerts the operator.[25]

Gating on uncertainty is necessary but not sufficient, because it works only where the model recognises that it is unsure; an input unlike anything in its training data, or one deliberately crafted to deceive, can produce a confident error that never trips the threshold, and a fixed floor such as 85 per cent carries the same weakness. It must therefore be paired with checks that do not depend on the model's own self-assessment: flagging inputs that fall outside the distribution on which the model was trained, corroborating a classification against a second model or an independent sensor, and testing the deployed suite against deliberate deception rather than only against representative signals. The third is explainability. A saliency map laid over the spectrogram, or the attention weights of a transformer, shows the operator which part of the signal drove the call.[26] This is not proof of why the model concluded, but where the model looked, and it serves as an aid to judgement rather than a substitute for it.

Autonomy and Its Limits

The objection to all of this is serious and should be stated in its strongest form. Deep reinforcement learning agents are probabilistic and lack formal verification, so a change of sea state or a domain shift could produce erratic jamming that disrupts friendly command links or escalates a situation nobody intended. A black-box failure cannot be diagnosed in real time, eroding situational awareness and deskilling the crew who must take over. And integrating AI merely trades conventional cyber vulnerabilities for adversarial ones, which a capable opponent can exploit to paralyse a multi-million-dollar platform. On this view, a deterministic rule-based system is slower but wholly predictable, verifiable and under human command, and therefore safer.

The rebuttal rests on three points. First, the rule-based alternative is not slow but blind: against a waveform not in the library, it offers nothing, and the certainty it provides is the certainty of not seeing the threat. Second, the choice is a false binary. Coupling deep-learning classification with deterministic safety boundaries and uncertainty gating produces a hybrid in which the AI operates inside a sandbox and surrenders control the moment confidence falls.

Third, the architecture elevates the human, rather than removing them completely. An operator freed from reading pulse trains becomes a supervisor of the machine's

[25] Suleiman G. Suleiman, A. E. Airoboman and Nathaniel Dalla, **"Artificial Intelligence-Driven Electronic Warfare Systems: The Future of Spectrum Dominance in Modern Conflict"**, no. 13.

[26] Nicolae Sfetcu, **"Challenges and Risks of Artificial Intelligence in Electronic Warfare"**, *IT & C*, Vol. 4, No. 1, 2025.

reasoning rather than a component in its processing chain. This is consistent with the sub-second window described earlier only because the two refer to different decisions. Responses inside the engagement window, such as countermeasure selection and release, are automatic and bounded by the rules the operator has set in advance. Everything slower—changes to those rules, the authorisation of a new technique, the review of what the suite did and why—waits for the operator. The human supervises by setting the boundary before the engagement and auditing the record afterwards, rather than by intervening during it. The safeguard is not the existence of the mechanism but the discipline of its use, and that is a question of doctrine and training as much as of software.

Recommendations

At the tactical level, existing naval platforms should be retrofitted with heterogeneous edge accelerators such as neuromorphic processors or low-power tensor processing units that can run wideband convolutional and recurrent models locally within SWaP-C and thermal limits. Onboard, EW subsystems should adopt open-architecture interfaces such as the NATO Generic Vehicle Architecture so sensing, classification, and execution modules can exchange data at low latency. Uncertainty gating should be built into the CMS as a hard interlock, blocking autonomous jamming or chaff release and handing control to the operator whenever classification confidence falls below a defined floor. A starting point such as 85 per cent, to be calibrated through trials.

At the operational level, UAV swarms carrying lightweight software-defined cognitive payloads should be developed as forward-deployed sensing and cooperative jamming nodes for task groups, with endurance addressed through hydrogen fuel-cell propulsion. Self-healing mesh protocols should be standardised across uncrewed assets so that the local network survives jamming and node loss. No such suite should enter service without closed-loop, high-fidelity simulation that exposes the models to diverse adversarial waveforms and realistic sea-state noise.

At the strategic level, three investments matter more than any single platform. First is sovereign synthetic RF data generation, where labelled emissions from a real adversary are classified or unobtainable in peacetime. High-fidelity national electromagnetic simulators are the only way to train models before contact, and the simulator becomes a strategic asset in its own right. Second, India must establish a standing, independent RF red team, empowered to test deployed models for evasion, poisoning, and extraction both before fielding and continuously afterwards. Third, India needs a settled legal and policy position. Autonomous spectrum decisions need defined accountability, and India should shape international guidelines rather than

receive them. Domestically, a joint EW test and evaluation establishment should carry these strands through the indigenisation routes already in place, so the architecture is proven against Indian emitters, Indian sea states, and Indian rules of engagement rather than imported wholesale.

Conclusion

The contest for the EM spectrum has moved from the physics of jamming to the speed of decision. Legacy suites were built on the assumption that the threat could be catalogued before the fight began, and that assumption no longer holds. An edge-deployed, multi-agent architecture that senses with a digital array, reconstructs with compressive sensing and classifies with deep learning, decides with a domain language model, and attacks with reinforcement learning offers a way to close the gap between detecting an emitter and countering it. None of it is safe unless the suite can also be trusted: adversarial training, uncertainty gating and honest explainability are not additions to the architecture but part of it. The risk of fielding such a suite is real, but the certainty of meeting agile, software-defined threats with static libraries is worse.

About the Author

Cdr. Vishnu B. Nair is a Communication & Electronic Warfare specialist, presently undergoing the 33rd Naval Staff and Technical Management Course (NSTMC-33) at the Naval War College, Goa.

Manohar Parrikar Institute for Defence Studies and Analyses is a non-partisan, autonomous body dedicated to objective research and policy relevant studies on all aspects of defence and security. Its mission is to promote national and international security through the generation and dissemination of knowledge on defence and security-related issues.

Disclaimer: Views expressed in Manohar Parrikar IDSA's publications and on its website are those of the authors and do not necessarily reflect the views of the Manohar Parrikar IDSA or the Government of India.

© Manohar Parrikar Institute for Defence Studies and Analyses (MP-IDSA) 2026

Manohar Parrikar Institute for Defence Studies and Analyses
1, Development Enclave, Rao Tula Ram Marg
New Delhi 110 010 India
T +91-11-2671 7983 F +91-11-2615 4191
www.idsa.in
Twitter @IDSAIndia
www.facebook.com/ManoharParrikarInstituteforDefenceStudiesAnalyses