International Journal of Future Engineering Innovations

Brain-Inspired Cognitive Architecture for Non-Centralized Spectrum Intelligence

Ameer Sameer Hamood Mohammed Ali [1*], Ali Samir Saleem [2]
[1] Presidency of the University of Babylon, University of Babylon TOEFL Center, Babylon, Iraq
[2] Al-Furat Al-Awsat Technical University, Babylon, Iraq

* Corresponding Author: **Ameer Sameer Hamood Mohammed Ali**

Article Info

ISSN (online): 3049-1215
Impact Factor (RSIF): 8.25
Volume: 03
Issue: 05
Received: 16-06-2026
Accepted: 18-07-2026
Published: 20-08-2026
Page No: 01-09

Abstract
Objective: To develop and evaluate a brain-inspired cognitive architecture that enables non-centralized spectrum intelligence while preserving primary-user (PU) protection and local radio autonomy.
Methods: Each cognitive radio is modeled as an autonomous agent with perception and attention, working and long-term memory, contextual reasoning, adaptive goal management, learning, and decision execution. These functions are organized into a twelve-stage perception-memory-reasoning-action lifecycle with optional bounded peer exchange, a PU-safety gate, convergence monitoring, low-overhead steady-state operation, and event-triggered reactivation. The architecture is algorithm-agnostic and can host neuro-fuzzy, swarm, evolutionary, reinforcement-learning, or hybrid mechanisms.
Results: A proof-of-concept 15-agent network achieved 0.7340 accuracy, 0.8670 specificity, 0.1330 false-positive rate, and 0.7068 ROC AUC. Recall was 0.2569 and F1 score was 0.2963, indicating a conservative operating point that reliably rejects unsafe channels but misses many usable opportunities.
Conclusion: The proposed architecture demonstrates the feasibility of persistent, context-aware, non-centralized spectrum cognition. The main improvement priorities are confidence-threshold calibration, opportunity detection, larger-scale repeated experiments, and module-level ablation studies.

DOI: https://doi.org/10.54660/IJFEI.2026.3.5.01-09

Keywords: agentic wireless intelligence, brain-inspired computing, cognitive radio network, dynamic spectrum access, multi-agent systems, non-centralized intelligence

1. Introduction

The transition toward 6G and highly heterogeneous wireless systems is intensifying the mismatch between static spectrum assignment and time-varying spectrum demand. Cognitive radio networks (CRNs) address this mismatch by sensing, learning, and dynamically accessing underutilized spectrum while respecting incumbent primary-user (PU) constraints. Recent CRN research emphasizes that sensing reliability, adaptive resource allocation, and decision latency must be treated jointly rather than as isolated physical-layer tasks [1]. In parallel, artificial intelligence (AI) has become a central enabler for dynamic spectrum access because it can learn non-linear occupancy patterns, adapt to non-stationary traffic, and coordinate spectrum decisions under incomplete information [2].

The research frontier has therefore moved from single-node classification toward distributed intelligence. Multi-agent reinforcement learning (MARL) is increasingly studied for wireless networks in which many decision makers interact with a shared spectrum and must learn under partial observability [3]. Fully decentralized MARL has also been investigated explicitly for fair dynamic spectrum access, demonstrating the importance of eliminating dependence on a central scheduler at execution time [4]. At the sensing layer, recent work combines predictive modeling with wideband spectrum observation [5], quantization-aware deep learning at the edge [6], multiview graph neural networks (GNNs) [7], and GCN-LSTM spatio-temporal modeling

[8]. Cooperative CNN-LSTM models further show that spatially distributed user observations can improve the representation of occupancy dynamics [9].

Beyond conventional sensing, the concept of spectrum intelligence is becoming broader. Joint communication-and-sensing formulations increasingly treat subcarrier and beamforming decisions as coupled cognitive actions [10]. Spectrum cognition has also been reframed as semantic situation understanding rather than only energy detection [11], while foundation-model approaches such as SpectrumFM investigate reusable representations that can support multiple spectrum tasks [12]. Neuromorphic continual learning offers another route toward adaptive low-energy intelligence [13], and emerging agentic edge-network paradigms argue for autonomous agents that perceive context, reason over goals, and act with reduced dependence on central orchestration [14]. Recent decentralized spectrum-sharing work in heterogeneous 6G networks likewise reinforces the need for local coordination mechanisms that remain effective under diverse radio conditions [15].

1.1. Problem Statement and Research Gap

Despite these advances, most existing CRN solutions still optimize a limited portion of the cognitive cycle. A spectrum-sensing model may classify a channel accurately but does not necessarily maintain persistent experiential memory, manage competing goals, explain when to become conservative, or decide how to re-enter learning after an environmental change. Similarly, distributed MARL can provide local action policies, yet the policy itself is not a complete cognitive architecture: perception, attention, working memory, long-term memory, reasoning, goal management, safe execution, and event-triggered re-adaptation are often left implicit or implemented as task-specific components. This fragmentation makes it difficult to transfer intelligence across spectrum tasks and creates engineering dependence on particular learning algorithms.

The central problem addressed in this paper is therefore: how can each cognitive radio be endowed with a reusable, brain-inspired perception-memory-reasoning-action loop that supports autonomous local decisions, bounded cooperation with peers, persistent learning, and PU-protection constraints without requiring a mandatory central controller? The challenge is not only algorithmic accuracy. The architecture must also remain scalable, maintain safe behavior under uncertainty, preserve useful experience over time, and react selectively to environmental events rather than continuously executing a computationally expensive global optimization.

1.2. Research Aim

The aim is to design and evaluate a brain-inspired cognitive architecture in which every secondary-user radio operates as an autonomous cognitive agent. The architecture maps biological cognition concepts to practical CRN functions: sensory perception to spectrum sensing, attention to channel/context prioritization, working memory to short-horizon state retention, long-term memory to experiential knowledge, reasoning to multi-objective action selection, neuroplasticity to adaptive learning, and motor action to spectrum-access execution.

1.3. Contributions and Novelty

The main contributions are as follows:

- A unified cognitive-agent model that integrates perception, attention, working memory, long-term memory, reasoning, goal management, learning, and action execution within each radio rather than treating sensing or resource allocation as an isolated AI task.
- A non-centralized network topology in which local agents can make spectrum-access decisions without a mandatory global controller; peer knowledge exchange is optional and bounded, reducing dependence on a single coordination point.
- A twelve-stage closed-loop methodology covering initialization, perception, knowledge representation, bio-inspired decision formation, utility evaluation, adaptation, spectrum access, transmission, feedback, convergence, steady-state cognition, and event-triggered reactivation.
- An algorithm-agnostic decision-and-learning interface that can accommodate neuro-fuzzy rules, swarm mechanisms, evolutionary adaptation, reinforcement learning, or hybrid combinations while preserving a common cognitive architecture.
- An explicit safety gate that makes PU protection and interference constraints admissibility conditions for action execution, including a defer/idle action when confidence or constraints are unsatisfactory.
- A proof-of-concept evaluation using 15 autonomous cognitive radio agents and a broad set of discrimination, calibration, agreement, and error metrics, enabling a more informative diagnosis than accuracy alone.

The novelty lies in the architectural integration of persistent memory, context-sensitive reasoning, adaptive goals, and non-centralized execution. Rather than proposing another stand-alone detector or another single optimization policy, the framework specifies how diverse learning mechanisms can coexist inside a reusable cognitive lifecycle. In this sense, the proposed system is positioned between low-level spectrum sensing and high-level agentic wireless intelligence: it provides a concrete radio-level architecture for sustained autonomous cognition.

2. Related Works

Recent 2026 literature can be organized into three closely related directions: intelligent spectrum sensing, distributed/federated multi-agent spectrum management, and emerging agentic or neuromorphic wireless intelligence. These directions provide important building blocks, but they differ in the extent to which they implement persistent cognition at each radio.

2.1. Intelligent and Collaborative Spectrum Sensing

Yi and Liang survey collaborative spectrum sensing from an AI perspective and organize learning-enabled approaches for combining distributed observations [16]. Ahmed *et al*. propose a multi-branch concatenated neural architecture for intelligent spectrum sensing [17], emphasizing richer learned representations at the detector. Reed and Sangwan combine physics-based machine-learning models with large-language-model reasoning at the edge for spectrum sensing [18], indicating a shift from pure classification toward reasoning-assisted interpretation. These works strengthen perception, but perception alone does not specify how a radio should maintain long-term experience, balance multiple goals, execute safe access decisions, and re-adapt after environmental events.

2.2. Distributed, Federated, and Multi-Agent Spectrum Management

Dhuheir *et al*. investigate digital-twin-assisted adaptive multi-agent deep reinforcement learning for spectrum and resource management in Open-RAN UAV-enabled 6G networks [19]. Khaf and Kaddoum develop federated hierarchical reinforcement learning for resilient spectrum sharing in non-terrestrial networks [20], while Batool *et al*. embed privacy considerations into attention-driven multi-agent spectrum coordination for 6G V2X networks [21]. Tashman and Cherkaoui address trustworthy AI control and reward-poisoning resilience for joint optimization in cognitive networks with dynamic hybrid reconfigurable intelligent surfaces [22]. Collectively, these studies demonstrate the value of distributed learning, federation, hierarchy, attention, and robustness. Their primary contribution, however, is an optimization or learning mechanism for a defined task rather than a reusable cognitive architecture that explicitly connects perception, dual-timescale memory, reasoning, goal management, execution, and event-triggered lifecycle control.

2.3. Agentic and Neuromorphic Wireless Intelligence

Agentic AI has begun to influence the design vocabulary of 6G networks. Ferrag *et al*. argue for agentic AI-native networks in which autonomous agents become persistent components of the network intelligence plane [23]. Li *et al*. propose intent-aware and continuously evolving physical-layer intelligence as part of an agentic wireless communication vision [24]. Semerikov *et al*. study neuromorphic computing and hardware-software co-design for energy-efficient cognitive radio spectrum intelligence [25]. These works motivate autonomy, intent, continual adaptation, and brain-inspired efficiency. The proposed architecture complements them by specifying a concrete, radio-level cognition loop for spectrum access and by making the division among perception, memory, reasoning, learning, action, and event-driven reactivation explicit.

2.4. Positioning of the Proposed System

Table 1 shows that recent work is rapidly advancing individual ingredients needed for spectrum intelligence, including collaborative sensing, distributed RL, privacy, robustness, agentic reasoning, and neuromorphic efficiency. The proposed system differs in four structural respects. First, it treats cognition as a complete local lifecycle rather than a single learned mapping. Second, it separates working memory from long-term experiential memory so that immediate context and accumulated experience can influence decisions differently. Third, it exposes goal management and PU-protection constraints as explicit reasoning inputs rather than embedding all preferences implicitly in a training reward. Fourth, it supports low-overhead steady-state operation with event-triggered reactivation, allowing the expensive cognition loop to be engaged when spectrum conditions change.

The comparison can be summarized by four architectural differentiators:

- Complete perception-memory-reasoning-action lifecycle at every radio.
- Explicit working-memory and long-term experiential-memory separation.
- Local goal management with PU-protection constraints and a safe defer/idle action.
- Event-triggered reactivation and optional bounded cooperation without a mandatory global controller.

Table 1: Comparison of recent 2026 related work and the proposed architecture

Ref.	Focus	Intelligence / coordination	Principal capability	Gap relative to the proposed brain-inspired architecture
[16]	Collaborative spectrum sensing survey	AI-enabled cooperative sensing	Taxonomy and synthesis of collaborative sensing methods	Survey/perception focus; does not define a persistent per-radio perception-memory-reasoning-action lifecycle.
[17]	Intelligent spectrum sensing	Multi-branch concatenated neural model	Rich learned detector representation	Detector-centric; no explicit long-term memory, goal manager, distributed cognitive lifecycle, or action feedback loop.
[18]	Edge spectrum sensing	Physics-based ML with LLM reasoning	Reasoning-assisted sensing at the edge	Reasoning augments perception, but spectrum-access execution, persistent radio memory, and multi-agent lifecycle are not the main focus.
[19]	UAV/Open-RAN spectrum and resource management	Digital twin + adaptive multi-agent DRL	Adaptive multi-agent resource decisions	Strong task-level learning; cognitive modules such as working/long-term memory and explicit event-triggered reactivation are not architecturally separated.
[20]	6G non-terrestrial spectrum sharing	Federated hierarchical RL	Resilient hierarchical spectrum sharing	Federated/hierarchical optimization is central; no unified brain-inspired agent architecture spanning sensing through action and memory.
[21]	6G V2X spectrum coordination	Attention-driven MARL with privacy mechanisms	Privacy-aware multi-agent coordination	Attention and coordination are explicit, but dual-timescale experiential memory and reusable cognitive lifecycle are not the organizing abstraction.
[22]	Cognitive MISO / dynamic hybrid RIS control	Trustworthy AI, joint optimization, poisoning resilience	Robust resource/control optimization	Focuses on robust RIS/network control rather than end-to-end cognitive radio autonomy and persistent spectrum memory.
[23]	AI-native 6G networking	Agentic AI vision	Autonomous network agents and orchestration	Broad network-level vision; the proposed work instantiates a concrete spectrum-access cognition loop at each radio.
[24]	Agentic physical-layer intelligence	Intent-aware continuously evolving agents	Intent-driven PHY adaptation	PHY/intent focus; the proposed system adds explicit spectrum perception, memory hierarchy, decision safety, and event-triggered steady-state operation.
[25]	Neuromorphic cognitive radio	Neuromorphic HW/SW co-design	Energy-efficient low-latency spectrum intelligence	Hardware/neuromorphic emphasis; the proposed framework addresses distributed cognitive-agent organization and algorithm-agnostic reasoning/learning.

3. Proposed Methodology

The proposed methodology converts each cognitive radio into a persistent cognitive agent. The radio does not wait for a centralized controller to provide a final action. Instead, it maintains a local internal state, senses its environment, selectively attends to relevant observations, retrieves short- and long-term experience, reasons over candidate actions and constraints, executes an admissible decision, and learns from the outcome. Cooperation is permitted through compact neighbor messages, but global coordination is not a prerequisite for operation.

3.1. Cognitive-Agent State Model

For radio i at decision time t, the local observation vector can be represented as a normalized state containing received signal quality, channel quality, interference, estimated PU occupancy, residual energy, and traffic/QoS context:

$$x_i(t) = [SNR_i(t), Q_i(t), I_i(t), P_{PU,i}(t), E_i(t), T_i(t)].$$

This state enters working memory $W_i(t)$, which maintains short-horizon context required for temporal reasoning. Long-term memory L_i stores compact experience tuples describing context, action, reward/utility, constraint violations, and outcome. A general memory update can be written as:

$$W_i(t) = f_W(W_i(t-1), x_i(t), a_i(t-1), r_i(t-1));$$

$$L_i(t) = f_L(L_i(t-1), W_i(t), outcome_i(t)).$$

The architecture deliberately does not prescribe a single implementation of f_W or f_L. Working memory may be a bounded buffer, recurrent hidden state, or attention cache; long-term memory may be a compressed replay store, rule base, semantic prototype library, or learned parameter memory. The architectural requirement is that immediate context and accumulated experience remain distinguishable and accessible to the reasoning module.

3.2. Perception, Attention, and Knowledge Representation

Perception begins with local spectrum sensing and feature extraction. Each agent estimates energy or signal statistics, interference, channel quality, and PU activity. An attention mechanism then prioritizes channels, neighbors, and contextual variables that are most relevant to the current goal. This can be implemented by learned attention weights, heuristics, uncertainty scores, or salience thresholds. The selected information is encoded into a consistent internal representation such as normalized state variables, fuzzy linguistic values, learned embeddings, or fitness components. The knowledge representation stage provides the bridge between raw sensing and reasoning. It stores channel histories, interference patterns, and previous strategy outcomes. Depending on the selected learning mechanism, the representation may include pheromone values, fuzzy rules, neural parameters, experience buffers, or other compact traces. The key design constraint is that knowledge must be locally available so that a radio can continue operating even when a peer or coordinator becomes unavailable.

3.3. Reasoning, Utility Evaluation, and Safety

The reasoning module produces a set of candidate actions $A_i(t)$, where an action may include channel selection, transmit power, scheduling decision, cooperative mode, or deliberate deferral. The decision engine balances exploration of new opportunities against exploitation of successful past strategies. Candidate actions are scored using a normalized multi-objective utility:

$$U_i(a,t) = w_T Tput_i - w_D Delay_i - w_E Energy_i - w_I Interf_i + w_F Fair_i + w_R Rel_i,$$

where the non-negative weights express the active goal profile and the component terms are normalized to comparable scales. The utility is evaluated only after hard protection constraints are checked. In particular, if predicted interference to an incumbent PU exceeds an admissible threshold or channel-occupancy confidence is inadequate, the action is rejected. This two-level decision structure separates safety constraints from preference optimization. A conservative idle/defer action is always retained as a safe fallback.

3.4. Adaptation and Learning

After each action, the agent compares the observed outcome with its expected utility and updates its internal parameters. The update mechanism is intentionally modular. Swarm-based instantiations may update and evaporate pheromone traces or velocity/position variables; neuro-fuzzy instantiations may tune rules and memberships; neural models may update weights; reinforcement-learning agents may update value or policy estimates; and evolutionary mechanisms may modify strategy populations. Inefficient strategies are progressively forgotten or down-weighted, whereas strategies that satisfy performance and PU-protection objectives are reinforced.

This stage implements the architectural analogue of neuroplasticity. Crucially, learning is not limited to a one-time training phase. The radio retains lightweight online adaptation so that past outcomes can change future priorities, while bounded memory and event-triggered reactivation prevent unbounded computational and storage growth.

3.5. Non-Centralized Cooperation Model

Each radio is capable of completing the cognition loop using local sensing and memory. When cooperation is available, agents may exchange compact information such as channel confidence, selected action, local interference estimate, compressed experience score, or neighbor reliability. These messages influence reasoning but do not constitute a mandatory central fusion point. If a neighbor fails, becomes unreachable, or provides low-confidence information, the local agent can continue with reduced cooperative input. This topology reduces single-point failures and allows heterogeneous radios to adopt different learning mechanisms while sharing a common interaction interface.

3.6. Twelve-Stage Cognitive Lifecycle

Figure 1 updates the original workflow into a continuous event-driven cognitive loop. The twelve stages are described below.

1. Initialization. Deploy SUs and PUs, initialize cognitive agents, memory structures, learning parameters, channel availability, PU activity models, communication range, noise/interference conditions, and traffic context.

2. Spectrum environment perception. Sense the current radio environment; estimate signal/energy levels, interference, channel quality, and PU activity; allocate attention to the most relevant channels and observations.

3. Knowledge representation. Encode observations into normalized state, fuzzy/linguistic variables, learned embeddings, or fitness features; update channel history, interference memory, and recent experience.

4. Bio-inspired decision engine. Generate candidate actions through exploration and exploitation; optionally use cooperative peer information, swarm signals, learned rules, neural inference, or a hybrid mechanism.

5. Decision evaluation. Compute utility and verify QoS, throughput, delay, energy, fairness/reliability, and PU-interference constraints; reject inadmissible actions before execution.

6. Adaptation and learning. Update internal parameters from outcomes; reinforce successful strategies and decay ineffective strategies using the selected bio-inspired learning mechanism.

7. Spectrum access decision. Finalize the admissible channel, power, and schedule decision, or deliberately defer when safety/confidence conditions are not met.

8. Transmission execution. Execute data transmission and optional cooperative communication using the selected action.

9. Feedback collection. Observe throughput, packet loss, interference events, PU reappearance, failures, and realized utility; write the outcome to working/long-term memory.

10. Convergence check. Assess whether the local policy/strategy and network behavior are sufficiently stable. If not, return to perception and continue adaptation.

11. Steady-state cognitive operation. Maintain the current safe access pattern with low-overhead monitoring, bounded incremental learning, and energy-aware operation.

12. Event-trigger detection. Detect PU reappearance, interference spikes, QoS degradation, topology change, confidence loss, or other salient events; reactivate full cognition when a trigger occurs.

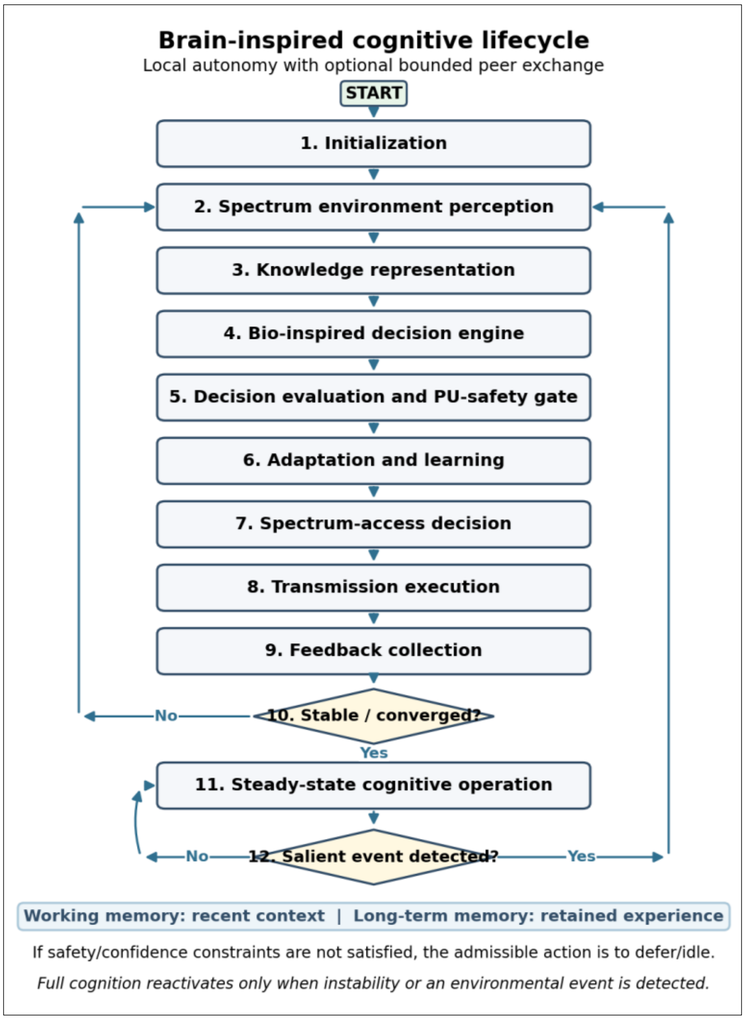

Fig 1: Updated flowchart of the proposed brain-inspired, non-centralized cognitive lifecycle.

3.7. Execution Logic

Algorithm 1: Local execution logic for cognitive radio i

Step	Operation
1	Initialize W_i, L_i, goal profile G_i, thresholds, and learning parameters.
2	Sense local spectrum and construct x_i(t); prioritize observations using attention/salience.
3	Retrieve relevant recent and long-term experience.
4	Generate candidate spectrum-access actions, including a safe idle/defer action.
5	Filter candidate actions using PU-protection and interference constraints.
6	Score admissible actions using U_i(a,t) and select the highest-utility action subject to the active goal profile.
7	Execute channel/power/schedule decision; exchange bounded peer information when enabled.
8	Observe outcome; update W_i, L_i, and learning parameters.
9	If behavior has not stabilized, return to Step 2; otherwise enter low-overhead monitoring.
10	If an event trigger occurs, return to Step 2 and reactivate the full cognition loop.

3.8. Scalability and Implementation Considerations

Because sensing, reasoning, and learning are local, per-agent computation scales primarily with the number of sensed channels, candidate actions, and the complexity of the selected learner rather than the total network size. Communication overhead depends on the number of neighbors with which compact cooperative information is exchanged. Bounded memories, neighbor selection, confidence-gated message use, and event-driven steady-state monitoring are therefore important implementation controls. For dense deployments, the architecture can restrict peer exchange to the k most informative neighbors or use periodic compressed summaries instead of broadcasting complete local state.

The architecture also supports graceful degradation. A radio can continue operating with local sensing if peer information is delayed or unavailable. Conversely, when sufficient connectivity exists, shared confidence and experience can accelerate adaptation. This separation between mandatory local cognition and optional cooperation is central to the non-centralized design.

4. Results

4.1. Proof-of-Concept Network Behavior

The available proof-of-concept implementation contains 15 cognitive radio agents. The connectivity snapshot in Figure 2(a) illustrates a dense peer-interaction graph, while Figure 2(b) shows node mobility trajectories across the simulated coordinate space. Together, the plots demonstrate that the architecture was exercised under changing spatial relationships rather than only a static single-link configuration. Because the intelligence is held at the agents, mobility changes neighborhood information without changing the underlying cognitive pipeline.

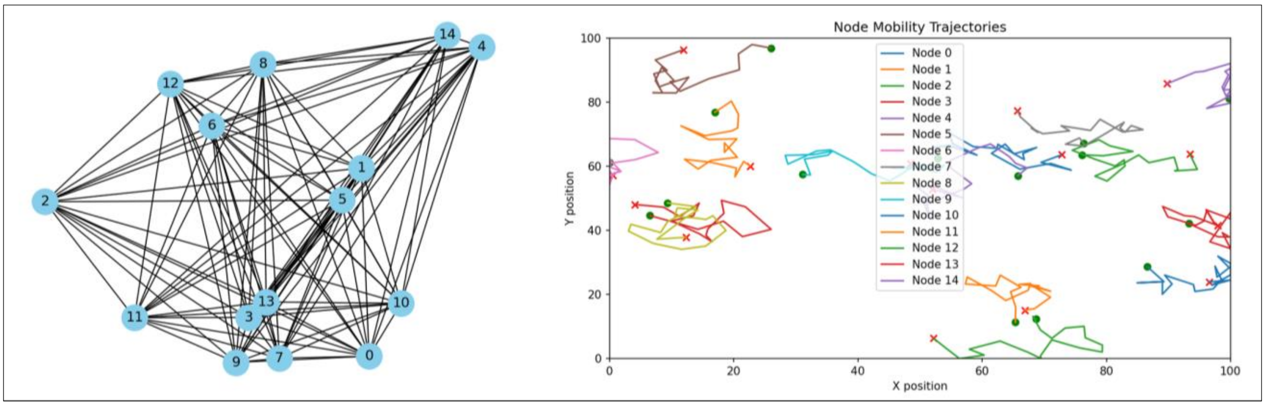

Fig 2: Proof-of-concept network outputs: (a) connectivity snapshot and (b) mobility trajectories for the 15-node cognitive-radio scenario.

4.2. Decision-Performance Metrics

Table 2 reports the complete set of metrics produced by the brain-inspired cognitive radio experiment. The metrics should be interpreted jointly because the operating point is asymmetric: false access to an occupied/unsafe channel is more damaging than missing a possible access opportunity. Under this setting, raw accuracy alone can obscure whether the model favors the safe or opportunity-seeking class.

Table 2: Reported performance metrics for the brain-inspired cognitive radio network

Metric	Value	Metric	Value
Accuracy	0.7340	Precision	0.3500
Recall (sensitivity)	0.2569	F1 score	0.2963
Specificity	0.8670	False positive rate	0.1330
False negative rate	0.7431	Negative predictive value	0.8071
Matthews correlation coefficient	0.1395	Balanced accuracy	0.5619
Jaccard index	0.1739	Cohen kappa	0.1370
Hamming loss	0.2660	Zero-one loss	0.2660
ROC AUC	0.7068	Average precision	0.3762
Log loss	0.4887	Brier score	0.1587
Detection rate	0.2569	False alarm rate	0.1330
Cognitive radio agents (SUs)	15	Total devices/nodes	15

The system reaches an overall accuracy of 0.7340, but the more diagnostic pattern is the combination of high specificity (0.8670) and low recall (0.2569). A false-positive rate of 0.1330 means that unsafe/negative cases are rejected relatively reliably, whereas the false-negative rate of 0.7431 shows that many positive spectrum opportunities are not selected. Precision is 0.3500 and the resulting F1 score is 0.2963. This operating point is therefore conservative: it protects against inappropriate access more effectively than it captures all usable opportunities.

Discrimination is moderate rather than random. ROC AUC is 0.7068, indicating that the scoring mechanism contains useful ordering information even though the fixed decision operating point produces weak recall. Balanced accuracy (0.5619), Matthews correlation coefficient (0.1395), Cohen kappa (0.1370), and Jaccard index (0.1739) confirm that class-sensitive agreement remains limited. These metrics are more informative than accuracy when the class distribution and asymmetric consequences of errors matter.

Calibration-oriented measures also reveal room for refinement. The Brier score is 0.1587 and log loss is 0.4887, while average precision is 0.3762. Taken together with the ROC AUC, these results suggest that the architecture already produces a non-trivial confidence signal; however, the policy that converts confidence into spectrum access is not yet positioned to recover a high fraction of available opportunities. Adaptive thresholds, context-conditioned confidence calibration, and utility weights that explicitly balance PU protection against missed opportunities are natural improvements.

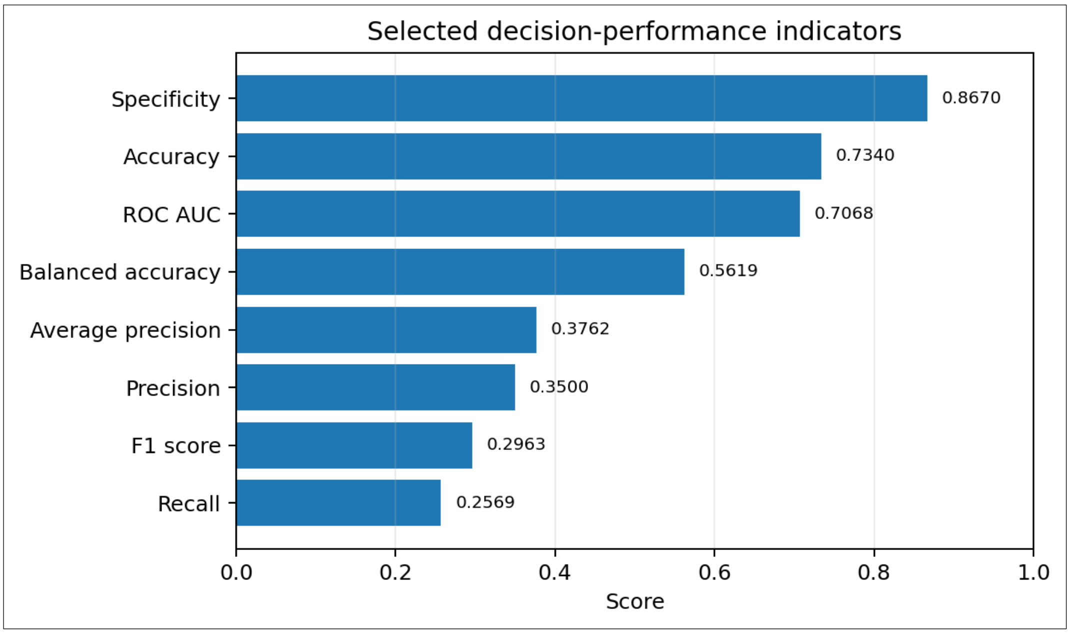

Fig 3: Visualization of selected reported decision metrics. High specificity contrasts with low recall and F1 score, revealing a safety-oriented operating point.

4.3. Diagnostic Results by Cognitive Function

Table 3: Diagnostic interpretation of the observed metrics

Cognitive / evaluation aspect	Evidence	Interpretation
PU-protection / safe rejection	Specificity 0.8670; FPR 0.1330; NPV 0.8071	Strongest aspect of the current operating point. The agent is comparatively reliable when rejecting unsafe/occupied cases.
Opportunity capture	Recall 0.2569; precision 0.3500; F1 0.2963; FNR 0.7431	Primary weakness. Many potentially usable opportunities are missed; threshold and reward/utility calibration should be targeted.
Ranking / discrimination	ROC AUC 0.7068; balanced accuracy 0.5619	The underlying score has useful discrimination, suggesting policy calibration may improve operating-point performance without redesigning every module.
Agreement and overlap	MCC 0.1395; kappa 0.1370; Jaccard 0.1739	Agreement beyond the dominant class is limited, consistent with conservative positive decisions.
Probability calibration	Brier 0.1587; log loss 0.4887; average precision 0.3762	Probabilistic confidence is usable but should be calibrated under changing PU activity and traffic conditions.
Network operation	15 autonomous agents with mobility/connectivity changes	The architecture can maintain local cognition across changing neighborhoods; larger-scale tests are needed to characterize overhead and convergence.

5. Discussion

5.1. Architectural Interpretation and Implications

The results support three conclusions about the proposed architecture. First, the perception-memory-reasoning-action pipeline can be instantiated as a functioning multi-agent spectrum decision system rather than remaining a conceptual diagram. Second, the system exhibits a clear safety bias: high specificity and low false-alarm rate are consistent with conservative reasoning under PU-protection constraints. Third, the gap between ROC AUC and recall indicates that future improvement should focus on the decision policy and confidence calibration in addition to the sensing representation itself. In practical terms, the architecture appears to know more than it currently acts upon: its ranking ability is moderate, but its final access threshold is restrictive. This distinction is important for a brain-inspired architecture. A cognitive agent should not be judged only by a single classifier score; the mapping from perception to action is conditioned by memory, goals, utility weights, constraints, and confidence. The present results therefore provide a useful diagnostic baseline for module-level ablation studies. For example, removing long-term memory can quantify the benefit of retained experience, disabling peer information can isolate cooperation gains, and replacing the adaptive goal profile with fixed weights can measure the value of context-dependent reasoning.

5.2. Limitations and Experimental Priorities

The current evaluation is a proof of concept and should not be interpreted as a claim of state-of-the-art superiority. The 15-agent scenario demonstrates autonomous multi-node operation, but larger and statistically repeated experiments are necessary to establish scalability and robustness. In particular, future evaluation should report confidence intervals over multiple random seeds, throughput/latency/energy measurements, convergence time, signaling overhead, PU collision probability, fairness, performance under sensing errors and jamming, and sensitivity to node density and mobility.

A rigorous ablation program should compare: (i) working memory only versus working plus long-term memory; (ii) local-only reasoning versus bounded cooperation; (iii) fixed versus event-triggered learning; (iv) fixed versus context-adaptive utility weights; and (v) alternative learners instantiated within the same architecture. Classical energy detection, supervised deep sensing, centralized/federated learning, and decentralized MARL baselines should be evaluated under identical channel traces. Such experiments would determine whether observed gains arise from the cognitive architecture itself rather than from a particular embedded learner.

6. Conclusion

This paper presented a brain-inspired cognitive architecture for non-centralized spectrum intelligence. The central design choice is to place a complete cognitive loop inside each radio: perception and attention interpret the local environment; working and long-term memory preserve context and experience; reasoning balances goals, utility, and PU-protection constraints; adaptation modifies behavior from feedback; and action execution closes the loop. A twelve-stage lifecycle adds convergence control, low-overhead steady-state operation, and event-triggered reactivation so that cognition remains persistent without requiring constant global optimization.

The proof-of-concept 15-agent results demonstrate the feasibility of this architecture and reveal a distinct operating profile. Accuracy is 0.7340, specificity is 0.8670, false-positive rate is 0.1330, and ROC AUC is 0.7068, indicating useful discrimination and comparatively strong rejection of unsafe cases. Recall of 0.2569 and F1 of 0.2963 show that the present policy is conservative and misses many usable opportunities. The immediate research priority is therefore not simply a more complex detector, but better calibration of how confidence, memory, goals, and constraints are converted into access actions.

Future work will evaluate the architecture at larger scale and under repeated channel realizations, quantify communication/energy overhead, and perform module-level ablations. Additional priorities include adversarial and jamming resilience, neuromorphic or edge-hardware implementations, adaptive confidence thresholds, memory consolidation, federated knowledge exchange without mandatory centralized inference, and comparison with contemporary decentralized MARL and agentic wireless baselines. These directions can transform the current architecture from a proof-of-concept neuro-cognitive radio model into a reproducible platform for autonomous spectrum intelligence in dynamic 6G environments.

References

1. Xu Y, Zheng K, Liu X, Li Z, Liu J. Cognitive radio networks: technologies, challenges and applications. Sensors. 2025;25(4):1011. doi:10.3390/s25041011.
2. Gbenga-Ilori A, Imoize AL, Noor K, Adebolu-Ololade PO. Artificial intelligence empowering dynamic spectrum access in advanced wireless communications: a comprehensive overview. AI. 2025;6(6):126. doi:10.3390/ai6060126.
3. Zhang J, *et al*. Multi-agent reinforcement learning in wireless distributed networks for 6G. arXiv. 2025;arXiv:2502.05812.
4. Zhang Y, Botelho P, Gordon T, Zussman G, Kadota I. Fair dynamic spectrum access via fully decentralized multi-agent reinforcement learning. arXiv. 2025;arXiv:2503.24296. To appear in: Proc WiOpt 2025.
5. Rojas A, Follet G, Jovanovic-Dolecek G, de la Rosa JM, Liñán Cembrano G. Prediction-based spectrum sensing framework for cognitive radio. IEEE Open J Circuits Syst. 2025;6:313-328. doi:10.1109/OJCAS.2025.3592376.
6. Abushahla HA, Varam D, AlHajri MI. Cognitive radio spectrum sensing on the edge: a quantization-aware deep learning approach. IEEE Commun Lett. 2025;29(8):1988-1992. doi:10.1109/LCOMM.2025.3582680.
7. Dong Y, Zhang M, Huo Y, Cheng X, Gao C, Wang J, *et al*. Multiview graph neural networks for spectrum sensing in cognitive radio. IEEE Internet Things J. 2025;12(15):29495-29513. doi:10.1109/JIOT.2025.3569323.
8. Wu Y, Liu Y, Wang F, Zhang Q, Wei G. Exploring spatio-temporal dynamics for spectrum sensing: a GCN-LSTM approach. IEEE Trans Veh Technol. 2025;74(12):19789-19794. doi:10.1109/TVT.2025.3583733.
9. Wang K, Chen Y, Bo D, Wang S. A novel multi-user collaborative cognitive radio spectrum sensing model: based on a CNN-LSTM model. PLoS One. 2025;20(1):e0316291. doi:10.1371/journal.pone.0316291.
10. Galappaththige D, Tellambura C. Wideband cognitive radio for joint communication and sensing: optimization of subcarrier allocation and beamforming. arXiv. 2025;arXiv:2503.17601.
11. Zhang H, Zhou F, Wu Q, Yuen C. Spectrum cognition: semantic situation for next-generation spectrum management. IEEE Netw. 2025. doi:10.1109/MNET.2025.3604901.
12. Liu C, *et al*. SpectrumFM: redefining spectrum cognition via foundation modeling. In: Proceedings of IEEE GLOBECOM. 2025. p.3861-3866. doi:10.1109/GLOBECOM59602.2025.11432674.
13. Liu Y, Qin Z, Zhu Y, Li GY. Enabling green wireless communications with neuromorphic continual learning. arXiv. 2025;arXiv:2502.17168.
14. Feng C, Zhang A, Min G, Huang Y, Quek TQS, You X. Towards 6G native-AI edge networks: a semantic-aware and agentic intelligence paradigm. arXiv. 2025;arXiv:2512.04405.
15. Tao Y, He JC, Liu ZJ, Yang S. Learn-to-share: a decentralized multi-agent spectrum sharing framework for heterogeneous networks in the 6G era. IEEE J Sel Areas Commun. 2026;44:3490-3506. doi:10.1109/JSAC.2026.3659811.
16. Yi P, Liang YC. Collaborative spectrum sensing in cognitive and intelligent wireless networks: an artificial intelligence perspective. arXiv. 2026;arXiv:2602.09615.
17. Ahmed T, Haque M, Sugiura Y, Shimamura T. A multi-branch concatenate neural approach for intelligent spectrum sensing in cognitive radio networks. IEEE Open J Comput Soc. 2026. doi:10.1109/OJCS.2026.3686814.
18. Reed A, Sangwan A. Edge intelligence using physics-based ML models with LLM reasoning for spectrum sensing. In: Proceedings of the 2026 International Conference on Smart Applications, Communications and Networking (SmartNets). 2026. doi:10.1109/SmartNets69662.2026.11604917.
19. Dhuheir M, Vu TX, Chatzinotas S. Digital twin-assisted adaptive multi-agent DRL for intelligent spectrum and resource management in Open-RAN UAV-enabled 6G networks. In: Proceedings of IEEE ICC. 2026. arXiv:2606.01324.
20. Khaf S, Kaddoum G. Federated hierarchical reinforcement learning for resilient spectrum sharing in 6G non-terrestrial networks. IEEE Open J Commun Soc. 2026. doi:10.1109/OJCOMS.2026.3693148.
21. Batool I, Fouda M, Ismail M, Verma S, Md Fadlullah Z, Kato N. Privacy-embedded spectrum coordination in 6G V2X networks: an attention-driven multi-agent learning approach. IEEE Netw. 2026. doi:10.1109/MNET.2026.3690057.
22. Tashman DH, Cherkaoui S. Trustworthy AI-driven dynamic hybrid RIS: joint optimization and reward poisoning-resilient control in cognitive MISO networks. IEEE Trans Netw Serv Manag. 2026. doi:10.1109/TNSM.2026.3660728.
23. Ferrag MA, Lakas A, Debbah M. 6G needs agents: toward agentic AI-native networks for autonomous intelligence. arXiv. 2026;arXiv:2605.01546.
24. Li Z, Jin X, Pan J, Yang Q, Shi Z. Agentic wireless communication for 6G: intent-aware and continuously evolving physical-layer intelligence. arXiv. 2026;arXiv:2602.17096.
25. Semerikov SO, Nechypurenko PP, Vakaliuk TA, Mintii IS, Kolhatin AO. Energy-efficient neuromorphic computing for ultra-low latency cognitive radio: a hardware-software co-design framework for 6G spectrum intelligence. Discover Artif Intell. 2026;6:363. doi:10.1007/s44163-026-01093-7.

How to Cite This Article

Mohammed ASHM, Saleem AS. Brain-inspired cognitive architecture for non-centralized spectrum intelligence. International Journal of Future Engineering Innovations. 2026;3(5):1-9. doi:10.54660/IJFEI.2026.3.5.01-09.

Creative Commons (CC) License