AI-ENHANCED WIRELESS SPECTRUM MANAGEMENT SYSTEM

Project Reference No.: 49S_BE_6947

College : REVA University, Bengaluru
Branch : Department of Electronics and Communication Engineering
Guide(s) : Dr. Nataraj Urs H.D
Student(s) : Mr. Rohan Baiju
Mr. Kunal R M
Mr. Nidhin R
Ms. Afra Khanam

Keywords:

Cognitive Radio, Spectrum Sensing, Convolutional Neural Network, Reinforcement Learning, Dynamic Spectrum Access

Introduction:

The radio frequency spectrum is the foundation of all wireless communication — from mobile phones and Wi-Fi to IoT sensors and broadcast television. With the rapid expansion of 5G networks and billions of connected devices, demand for spectrum has grown dramatically. Yet measurement studies consistently show that licensed spectrum bands remain idle 50–80% of the time. The problem is not physical scarcity but outdated static allocation policies that reserve bands for licensed Primary Users (PUs) even when they are not transmitting.

Cognitive Radio (CR) technology addresses this by enabling Secondary Users (SUs) to sense the radio environment and opportunistically use idle spectrum without interfering with licensed users. The core challenge is doing this accurately and in real time — signal conditions change rapidly and wrong decisions cause harmful interference. This project proposes an AI-enhanced solution combining a Convolutional Neural Network (CNN) for signal identification with a Reinforcement Learning (RL) agent for intelligent channel allocation, integrated into a live full-stack system running at a 500 ms cycle.

Figure 1: CNN Training History — validation accuracy stabilises above 94% over 36 epochs, confirming the model generalises well across all six signal classes.

Objectives:

1. Design and train a 1D-CNN to classify six radio signal types (Wi-Fi, LTE, Bluetooth, FM Radio, TV Broadcast, Empty) from raw I/Q samples with accuracy exceeding 90%.
2. Build a custom Cognitive Radio simulation environment where a Reinforcement Learning agent learns optimal channel allocation strategies through trial and reward feedback.
3. Train a PPO-based RL agent to assign Secondary Users to free channels while strictly avoiding any interference with licensed Primary Users.
4. Integrate both AI models into a real-time system with allocation decisions updated every 500 ms, logged to a database, and broadcast live over WebSocket.
5. Build a full-stack web dashboard displaying live channel allocations, CNN confidence scores, reward history, and a scrollable RL decision log.
6. Design the system architecture for hardware-ready deployment on a HackRF One SDR and NVIDIA Jetson Nano edge device in the next phase.

Methodology:

The system operates as a two-stage AI pipeline.

Stage 1 — Signal Classification: Radio signals are represented as I/Q (in-phase and quadrature) samples, the standard format in Software-Defined Radio. Each snapshot is a (2×256) array — 2 channels (I and Q) with 256 time steps. A 1D-CNN processes this data through four convolutional blocks (Conv1d → Batch Normalisation → ReLU → MaxPool), followed by Global Average Pooling and two fully connected layers with 50% dropout. The output is a 6-class softmax probability. A confidence threshold of 0.60 ensures only high-confidence labels are passed downstream, preventing uncertain predictions from triggering incorrect allocation decisions.

Stage 2 — RL-based Allocation: A custom Gymnasium environment models the cognitive radio network with 3 Secondary Users, 2 Primary channels, and a common band. The 10-dimensional state vector encodes Primary User activity, SU locations, CNN signal types, and channel quality scores. The PPO agent chooses from 9 actions: move any SU to either primary channel, recall an SU to common, or hold current state. The reward function gives +2.0 for placing an SU on a free channel (plus a quality bonus), +0.5 for a proactive recall, and −2.0 for any interference event. The agent was trained for 200,000 simulation steps using stable-baselines3. The full pipeline runs in a FastAPI backend: I/Q generation → CNN inference → RL decision → SQLite log → WebSocket broadcast to a Next.js frontend.

Result and Conclusion:

The CNN achieved 94.4% overall test accuracy. LTE, Bluetooth, FM Radio, and TV Broadcast were classified with 100% per-class accuracy. Minor confusion between Wi-Fi and Empty channel (both appear noise-like at low SNR) has negligible impact on system operation. The RL agent's decision accuracy grew from ~30% in early episodes to a stable 94.1%, consistently exceeding the 90% target. Channel utilisation averaged 70–80% versus ~50% under static allocation, demonstrating a clear improvement in spectrum efficiency. The end-to-end pipeline runs comfortably within 500 ms: CNN inference under 5 ms, RL decision under 2 ms, and WebSocket broadcast under 20 ms. In conclusion, the CNN + PPO approach successfully delivers real-time, interference-free intelligent spectrum management on a standard laptop — no dedicated hardware required.

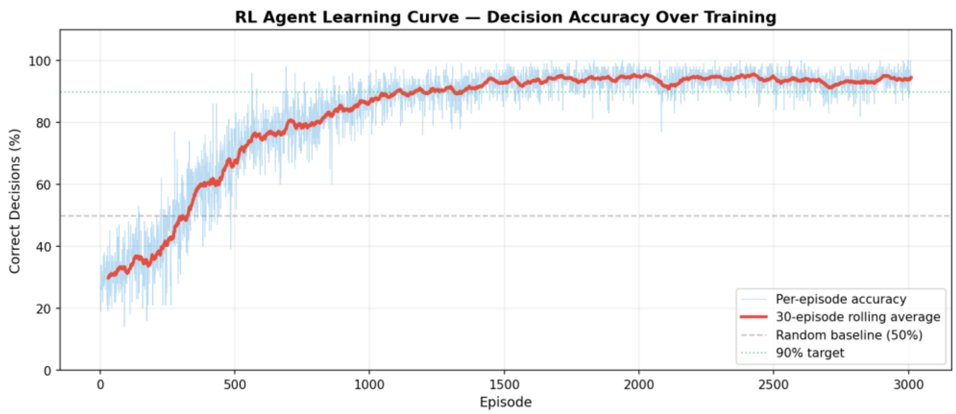

Figure 2: RL Agent Learning Curve — decision accuracy rises from ~30% to 94.1% over 3,000 training episodes, well above the 90% target (green line), showing genuine policy learning.

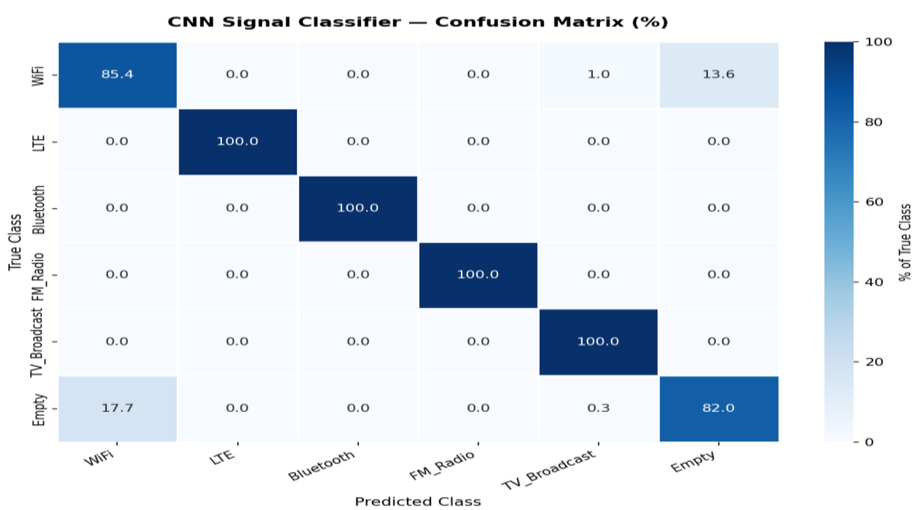

Figure 3: CNN Confusion Matrix — LTE, Bluetooth, FM Radio and TV Broadcast achieve 100% per-class accuracy. Minor confusion between Wi-Fi and Empty at low SNR is expected and does not impact system operation. This confirms the CNN reliably identifies signal occupancy before the RL agent makes allocation decisions.

References:

1. B. Sabir et al., "Systematic Literature Review of AI-enabled Spectrum Management in 6G," arXiv:2407.10981, 2024.
2. Upadhye et al., "A Survey on ML Algorithms for Cognitive Radio Networks," arXiv:2106.10413, 2021.
3. C. Cullen et al., "Predicting Dynamic Spectrum Allocation: A Review," Springer, 2023.
4. DARPA, "Spectrum Collaboration Challenge (SC2)," darpa.mil, 2016.
5. S. Haykin, "Cognitive Radio: Brain-Empowered Wireless Communications," IEEE JSAC, vol. 23, pp. 201–220, 2005.
6. J. Schulman et al., "Proximal Policy Optimization Algorithms," arXiv:1707.06347, 2017.
7. Raffin et al., "Stable-Baselines3: Reliable RL Implementations," JMLR, vol. 22, 2021.

Future Scope:

The future scope of this project includes:

1. Hardware integration: replace synthetic I/Q generation with a HackRF One SDR to capture real over-the-air signals, validating the CNN under actual noise and channel conditions.
2. Edge deployment of the entire AI pipeline on an NVIDIA Jetson Nano for low-power,real-time spectrum management at the device level.
3. Scaling the RL environment to more channels and Secondary Users, and exploring multi-agent RL where each SU independently learns its own allocation policy.
4. Improving CNN performance at low SNR using attention mechanisms or transformer-based architectures for more robust classification.
5. Federated spectrum management — multiple cognitive radio nodes at different locations collaboratively train a shared allocation policy without centralised data collection.