International Journal of
Human Computations & Intelligence
Vol. 04, Issue. 06, 2025

RESEARCH ARTICLE OPEN ACCESS

AI-Based Dynamic Spectrum Prediction and Allocation for IoT Wireless Networks Using Python

G Nagendra Prasad . K Riyazuddin

Department of Electronics and Communication Engineering,
Annamacharya University, Andhra Pradesh, India.

DOI: **10.5281/zenodo.17377265**
Received: 17 September 2025 / Revised: 11 October 2025 / Accepted: 17 October 2025

Abstract – The advent of the Internet of Things (IoT) has put pressure on scarce spectrum resources, especially in heterogeneous, interference-rich, and latency-critical environments. Static access regimes and rule-based ones cannot handle non-stationary interference and ultra-dense deployments. Leveraging advances in artificial intelligence (AI), this paper investigates and achieves spectrum intelligence: short-horizon spectrum prediction and dynamic, risk-aware allocation at radio timescales. We combine classical machine learning techniques, deep sequence and vision encoders (LSTM/GRU/TCN, spectrogram/REM models), transformers, reinforcement learning, and graph-based surrogates for channel–power assignment. Aside from modeling, reproducibility and deployability are also our focus with Pythonic pipelines: leakage-safe preprocessing, calibration (Brier/ECE), safety shields, and closed-loop evaluation with ns3/ns3-gym. Public datasets and simulators are inventoried with guidelines for splits and metrics (utilization, latency, violation rate, fairness, and energy). We cover engineering trade-offs for deployment on the edge (quantization, ONNX/TorchScript, federation with Flower) and outline open challenges in data sparsity, domain adaptation, explainability, security, and real-time feasibility. The result is an actionable roadmap to bring AI-enabled spectrum management from promising prototypes to robust, scalable IoT systems.

Index Terms – AI-based spectrum management; IoT wireless networks; dynamic spectrum access; spectrum forecasting; reinforcement learning; deep learning; graph neural networks; Python implementation; cognitive radio; 5G/6G.

I. INTRODUCTION

The radio spectrum is the foundation of the Internet of Things (IoT), powering everything from low-power periodic sensing to latency-sensitive control-plane signalling. With billions of devices that are

MILESTONE RESEARCH.IN OPEN ACCESS

ISSN (Online): 2583-5696
Int. Jr. of Hum Comp. & Int.

highly likely to run in heterogeneous environments, predicting when and where channels would be available has emerged as a pressing research issue. Early cognitive radio experiments demonstrated that simple machine learning (ML) baselines, trained from features such as duty cycle, frequency, and transmit power, could distinguish convincingly occupied versus free spectrum, establishing reproducible baselines for spectrum intelligence [1].

At the scale of smart cities, device diversity and bursty traffic, however, make demand highly volatile. Empirical evaluation on massive IoT traffic traces shows that good features and deep classifiers together can predict device behavior with extremely high accuracy [2]. This requires proactive spectrum prediction, where access policies are dynamically set downstream in anticipation of predictions of demand patterns in advance. But rule-based coexistence mechanisms and hard thresholds cannot realistically solve this issue. Experiments in illicit bands have shown that static approaches cannot handle non-stationary interference and stringent delay budgets, which present the necessity of acquiring learning-based control loops that can respond on radio timescales with coexistence guaranteed [3].

Recent advances in sequence learning enabled models that take advantage of temporal, spectral, and spatial dependencies within occupancy data. Composite recurrent forms, such as multi-dimensional LSTMs, provide principled approaches to encoding inter-step relationships for more precise multi-step forecasting at reasonable complexity [4]. Combined inter-channel relationships with temporal memory boost long-horizon prediction to allow anticipation-based reservation before bursts of contention [5]. Subsequent encoder-decoder transformer extensions with additional attention have since made prediction across multiple channels possible without the cost of prohibitive feature engineering, making it deployable in dynamic IoT environments [6].

Prediction alone is not sufficient; efficient allocation also has to be performed for handling the spectrum. Graph-neural surrogates have then emerged in the form of scalable near-optimal channel–power assignment approximators that balance accuracy and computation for dense IoT deployments [7]. Federated learning on Open RAN architectures presents an alternative path where distributed learning is performed without raw spectrum traces being transmitted out, hence eliminating privacy and non-stationarity concerns [8]. Reinforcement learning, particularly actor–critic methods, achieves explicit mapping from state space to spectrum allocation action, effectively closing the control–forecasting loop [9]. Other advancements, such as domain-knowledge-provided sequence models for geometric propagation smoothness representation [10] and distributed deep RL variants with good generalization under partial observability [11], bring us nearer robustness and practicability to real-time control. To this end, the focus of this paper is to investigate a general overview of AI-based approaches for spectrum prediction and allocation in IoT networks. Specifically, we will (i) compare the performance of multiple AI models in forecasting occupancy and demand in time, frequency, and space; (ii) elaborate how forecast data are integrated with adaptive allocation strategies; and (iii) introduce some recurring challenges such as the unavailability of datasets, explainability, and real-time applicability.

Our three-fold contributions are: (i) to show a single, integrating pipeline that spans predictive models to allocation policies with IoT constraints; (ii) to analyse datasets, metrics, and implementation practices, pointing out strengths and weaknesses; and (iii) to introduce open research challenges, such as data scarcity, explainability, and real-time feasibility. The guidelines outlined here fill the gap between theoretical AI models and implementable spectrum management systems for massive IoT deployments.

By building on prior cognitive radio intuition that merged prediction, sensing, and access [12], this paper builds an end-to-end spectrum intelligence overview to enhance large-scale IoT demands.

II. SYSTEMATIC LITERATURE REVIEW

Review scope and objectives

The review charts the state of the art at the intersection of AI-driven spectrum prediction and dynamic spectrum allocation for IoT-oriented wireless networks, with a practical emphasis on Pythonic implementation and replicability. The review spans 2013–2025, from pioneering cognitive-radio work on prediction to current advancements for 5G and prospective 6G systems. The chapter tracks three questions:

- which AI paradigms, traditional machine learning, deep learning, reinforcement learning, or hybrids, provide concrete gains for spectrum prediction and allocation,
- what are data regimes, metrics, and evaluation practices
- what obstacles still hinder deployment in heterogeneous, latency-sensitive IoT environments.

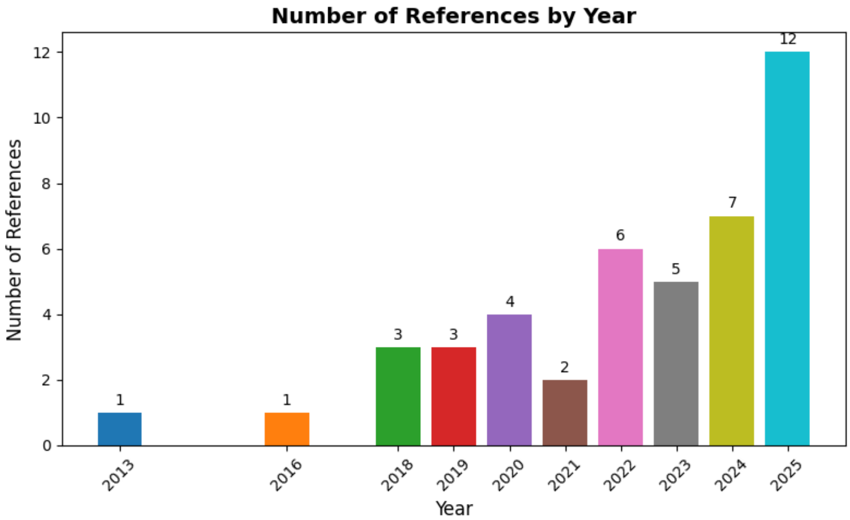

Fig. 1: Year-wise count of the papers

Identification

Systematic searching in IEEE Xplore, ScienceDirect, ACM Digital Library, SpringerLink, arXiv, and Google Scholar was done using the following combinations of keywords:

- "Spectrum Prediction" OR "Spectrum Occupancy Forecasting" OR "Radio Environment Map Prediction"
- "Dynamic Spectrum Access" OR "Spectrum Sharing" OR "Channel Allocation" OR "Power Control"
- "IoT" OR "Cognitive Radio Networks (CRN)" OR "5G" OR "6G"
- "Machine Learning" OR "Deep Learning" OR "Reinforcement Learning" OR "Graph Neural Network" OR "Hybrid Models"
- "Federated Learning" OR "Edge Intelligence" OR "Privacy-Preserving Spectrum Management"

Screening and eligibility

Screening occurred in two stages. Title–abstract screening retained papers that directly implied learning-based methods to spectrum prediction or dynamic allocation, or resources-optimization methods that were directly applicable to spectrum decisions (e.g., FK-LHSA, ADASA, RL-based DSS, hybrid RF–GRU–attention forecasting). Full-text analysis then confirmed the presence of an explicit algorithmic contribution and quantitative results (e.g., accuracy, throughput, latency, utilization, or energy).

Inclusion criteria

The final collection of primary studies was governed by the following criteria, modified to AI-based spectrum forecasting and dynamic assignment for IoT/CRN/5G/6G networks:

- Time window: Appeared in peer-reviewed journals or flagship conferences between 2018–2025.
- Problem focus: Focused spectrum forecasting (e.g., occupancy/demand/time-series forecasting) and/or dynamic selection/access/allocation of spectrum; tight cousin resource-optimization research was taken into account if the technique is tractably actionable for spectrum decisioning (e.g., FK-LHSA, ADASA, RL-based DSS).
- AI methodology: Employed machine learning, deep learning, reinforcement learning, or hybrid pipelines; algorithms described with sufficient procedural detail to enable Python implementation (model class, features/inputs, training setup).
- Quantitative evidence: Supplied task-relevant metrics, including at least one of:
- Prediction: Accuracy/AUC/MAE/MSE, or error over multi-step horizons;
- Allocation: Throughput, spectrum usage, collision/PU-violation rate, delay, fairness (e.g., Jain's index), or energy.
- Data transparency: Utilized public measurements, well-delineated simulations (e.g., NS-3/custom with traffic/interference models), or open RF datasets (e.g., Kaggle/WiSig), with the dataset/split clearly documented.

Fig. 2: Total count of the articles

Data extraction and thematic synthesis

The following fields were extracted from each included primary study to allow consistent analysis and reproducibility:

- Dataset features: measurement vs. simulation, public/private availability, bands and bandwidth, sampling/observation window, dataset size, and train/validation/test protocol.

- Model architecture: standard ML baselines (LR/SVM/RF/GBDT), temporal models (LSTM/GRU/TCN), Transformers for long-horizon prediction, CNN/vision backbones for spectrogram/REM inputs, RL (DQN/Actor–Critic/MARL), and graph-aware encoders (GNNs) when interference topology is represented.
- Representation & features: hand-crafted radio features (RSSI, duty-cycle, cyclostationary/radiometrics), spectrogram/REM embeddings, attention-based representations, graph features, and environment/context features (load, QoS, mobility).
- Training & evaluation setup: optimizer/schedule, horizon length for prediction, action space for allocation (channel/power/backoff), constraints (PU protection), and baselines/ablation definitions.
- Performance metrics: prediction error (Accuracy/AUC/MAE/MSE), throughput, utilization, collision/violation ratio, latency, fairness, energy, and convergence (for RL).
- Interpretability & diagnostics: feature importance (permutation/SHAP), saliency/attention maps on spectrograms/REMs, policy heat-maps and Q-value landscapes for RL, and failure-case analysis under distribution shift.
- Explicit limitations: scale or realism of the dataset, compute/latency overhead, generalization between bands/topologies, safety guarantees, reproducibility gaps.

The synthesis was organized around five core themes:

- Shift towards sequence and Transformer-based predictors for multi-step, multi-channel prediction (LSTM/GRU/TCN - Transformer hybrids).
- Simulation and data augmentation (NS-3/trace synthesis, domain adaptation from synthesized RF to real measurements) roles in data sparsity mitigation.
- Modality and representation constraints, IQ streams, spectrogram/REM, CSI, coupled with preprocessing to make them model-ready.
- Safety and explainability as first-class constraints, constrained or shielded RL and model attribution for auditability.
- Generalizability and deployment readiness, such as cross-band/site transfer, federated/edge learning, scalability for multi-agents, and standard reporting.

Reason and contribution of this review

A fragmented literature has grown in which prediction and allocation are typically solved as separate problems, evaluated on non-comparable data with heterogeneous measures and minimal reporting of failure modes or constraints. This survey adopts a pipeline-based view: prediction elements that forecast occupancy/demand are conditionally linked to allocation elements tha decide channel/power, so that progress in one stage is evaluated in terms of total utility (utilization, latency, fairness, and PU safety). The threefold contribution. First, a methods taxonomy is created that covers classical optimization, deep temporal/vision models, reinforcement learning, and graph-aware encoders. Within this taxonomy, the review specifies when each family is useful (e.g., Transformers for long-horizon predictions; MARL for decentralized distribution; GNNs when interference topology prevails) and what they cost (data gluttony, compute/latency, reproducibility sensitivity). Second, the review formalizes report standards in support of strict replication and comparison in Python: fixed seeds and splits; disclosure of observation window and action space; baselines mandated (e.g., optimization-only or heuristic); and a minimum set of metrics (prediction error + {throughput, utilization, collision/PU-violation, latency, fairness, energy}). The standards align with operational constraints within IoT settings and seek to reduce historical imprecision that has stunted cross-paper synthesis. Third, the review values credibility. For prediction, this entails stress-testing under distribution shift (band/site/time) and publishing attribution analysis. For allocation, particularly RL, it encourages constraint handling (hard shielding or constrained objectives), safety audits against incumbent tampering, and policy interpretability through state-action summaries. The survey also catalogs data streams, from public RF corpora (e.g., Kaggle/WiSig) to NS-3 simulation roll-outs, so researchers can pretrain at scale and subsequently validate on realistic small-scale traces.

Comprehensive Review

The rapid progress in wireless communication networks, such as cognitive radio networks (CRNs), 5G, and next generation 6G systems, has caused crucial challenges in spectrum utilization, resource allocation, network traffic forecasting, and optimal service quality, requiring intelligent, adaptive, and predictive techniques. Spectrum prediction as a central module of cognitive radio systems has been studied extensively to improve spectrum sensing, decision-making, sharing, and mobility. Xiaoshuang et al. [13] provided a detailed survey of prediction methods for the spectrum in CRNs, emphasizing the need for accurate prediction to reduce processing latency and improve spectrum efficiency, and identifying open problems such as dynamic channel availability, real-time adaptability, and prediction accuracy in heterogeneous environments. Building on these concepts.

Bharathi et al. [14] developed the AI-Assisted Adaptive Searching Algorithm (AIAS), which integrates supervised learning for predicting spectrum availability with reinforcement learning for real-time adaptive allocation to secondary users. Their method, tested on the Kaggle Spectrum Dataset, yielded 91.3% prediction accuracy, 92.5% average spectrum utilization efficiency, and interference-free allocations of up to 95.1%, cutting latency by about 50%, illustrating the real-world feasibility of AI in adaptive spectrum management in highly utilized environments. In the high-reliability and ultra-low latency of 5G networks that support the boom of IoT devices as well as ultra-dense deployments, Ramesh et al. [15] introduced the FK-LHSA scheme that combines Fractional Knapsack-based multi-band spectrum selection and Lagrange Hyperplane-based spectrum access allocation, optimizing the throughput, reducing spectrum access delay, and increasing the accuracy of allocation in IoT sensor networks. Adding to this, Banoth Ravi et al. [16] used machine learning algorithms to achieve intelligent, predictive, and adaptive spectrum allocation in 5G and beyond networks with real-time self-optimization and energy-efficient network resource management through accurate forecasting of spectrum demand. In cognitive communication paradigms for heterogeneous wireless networks, the application of AI has been shown to further enhance spectrum utilization. Kai Lin et al. [17] developed the AI-based Data Analytics-based Spectrum Allocation (ADASA) algorithm, integrating deep learning for feature extraction, dimensionality reduction, and user correlation analysis, allowing adaptive parameter tuning according to the network scenario and achieving significant gains in spectrum efficiency.

Similarly, Ramy et al. [18] proposed reinforcement learning-based Dynamic Spectrum Sharing (DSS) for next-generation networks, wherein the AI-bespoke framework selects resources dynamically based on existing network conditions and user activity, leading to enhanced LOS rates, throughput, spectrum efficiency, and latency minimization, while indicating continued challenges of data accessibility, algorithmic computational expense, security, and standardization. aside from spectrum management, Rawan et al. [19] explored the paradigm of personalized wireless networks with the introduction of an AI-driven, big data-powered, surrogate-assisted multi-objective optimization system that micro-manages scarce network resources at the individual user level with competing objectives of resource optimization and user satisfaction. The solution speaks to the feasibility of AI to provide real-time tailored service quality, opening smart network structures that move beyond stiff, preconceived QoS models. Network traffic prediction and QoS enhancement have also become critical in low-latency, high-speed environments. Ibrahim [20] addressed this by developing a stacking ensemble approach for network intrusion detection, backed by LSTM and LSTM Encoder-Decoder models for predicting 5G network QoS metrics like throughput, latency, jitter, and packet loss. Tested over a large-scale 50-day field dataset, the models achieved high attack detection rates (90.4% and 98.7%) with F1-scores of 90.0% and 98.5%, and low prediction errors for QoS metrics (14.57% and 13.75%), demonstrating AI's effectiveness for simultaneous security and service optimization in live 5G environments.

Scaling this to 6G networks, Mohammed Anis et al. [21] proposed a hybrid AI model of Random Forest (RF), combined with Gated Recurrent Units (GRU) and attention mechanisms, to accurately predict 6G network traffic in various channel conditions and user scenarios. The model achieved excellent performance metrics, including RMSE of 0.0049, MAE of 0.0034, MAPE of 0.46%, and R^2 of 0.9970, far surpassing baseline GRU and LSTM models, and enabling proactive resource planning, secure robustness, and real-time network optimization. Providing an overview, Karthick [22] surveyed AI methods applied in wireless communication such as machine learning, deep learning, and reinforcement learning to manage spectrum, allocate resources, detect signals, and predictive maintenance and identified open challenges like cross-layer integration, standardization gaps, computational overhead, and data privacy concerns, thereby giving a vision for the future of intelligent, self-organizing, and autonomous wireless networks.

Limitations of Existing Studies Identified

Despite the comprehensive development reported in employing AI for predicting the spectrum, allocating it, and optimizing wireless resources, several limitations have been identified in the existing studies, which restrict real-world implementation and scalability of these solutions.

1. **Absence of Real-World Validation:** Much of the work (e.g., Xiaoshuang et al. [13], Kai Lin et al. [17], Banoth Ravi et al. [16], Bharathi et al. [14]) is simulation- or controlled-dataset-based (e.g., the Kaggle Spectrum dataset). Simulations provide information on algorithmic performance but do not capture the full richness of dynamic, heterogeneous, interference-rich IoT and 5G/6G deployments. This poses a problem with the generalizability of the models described to real deployments.
2. **High Computational Complexity and Latency:** AI-based models such as ADASA [17], reinforcement learning-based DSS [18], and RF–GRU–Attention hybrid models [8] provide high accuracy but at the cost of computational load. These approaches will likely not scale for ultra-dense IoT environments where real-time processing with ultra-low latency is an absolute requirement. Similarly, ensemble and hybrid techniques require intense training, which may not be feasible on resource-constrained IoT devices.
3. **Scalability Issues in Dense IoT Networks:** FK-LHSA [15] and AIAS [14] schemes are effective in small-to-medium scale applications but are bound to fail under extreme IoT device connectivity.

Endless monitoring and re-allocation need create scalability bottlenecks that, accordingly, may cause delay, spectrum fragmentation, and increased interference under ultra-dense deployments.

4. **Large and High-Quality Dataset Dependence:** Various techniques (e.g., Rawan et al. [19], Ibrahim [20], Mohammed Anis et al. [21]) rely on big data and long-term datasets for training. In practice, it is expensive, time-consuming, and often impossible to acquire such datasets in dynamic spectrum environments. Furthermore, data sparsity, imbalance, and privacy concerns further limit the capability of data-driven AI models.
5. **Limited Cross-Layer and Cross-Network Integration:** Latest research aims at spectrum prediction or allocation on a single network layer only, without end-to-end integration between layers (physical, MAC, network) or among heterogeneous wireless systems (e.g., CRNs, 5G, 6G). According to Karthick [22], this does not allow for end-to-end optimization and hence creates gaps between spectrum management and overall network QoS, security, and energy efficiency.
6. **Security, Privacy, and Standardization Challenges:** Dynamic spectrum sharing frameworks (e.g., Ramy et al. [18]) also identify concerns of security threats, standardization deficiencies, and trust in AI-driven spectrum allocation. Since sensitive spectrum utilization information is processed by AI systems, concerns about adversarial attacks, data breaches, and the absence of standardized protocols for secure adoption persist.
7. **Energy and Resource Constraints in IoT Devices:** Whereas some articles (e.g., Banoth Ravi et al. [16]) focus on energy efficiency improvements, most of the existing frameworks assume high computational power and unlimited resources, which is unrealistic for low-power IoT devices. This incompatibility prevents the achievement of AI-based spectrum allocation in resource-constrained and battery-limited IoT settings.
8. **Restricted Investigation of Advanced Deep Learning Models:** Despite conventional ML and hybrid methods dominating existing research, advanced deep learning architectures (e.g., Transformers, Graph Neural Networks) remain underexploited for spectrum prediction and assignment. Bharathi et al. [14] also note that the integration of deep learning can possibly improve robustness and accuracy but is more a future research direction rather than an adopted method

Table 1: Overview of the existing literature

Reference	Model / Algorithm	Strength	Limitation
[13]	Survey of Spectrum Prediction Techniques	Comprehensive overview of spectrum sensing, decision, sharing, and mobility; identifies key challenges and research directions	No experimental validation; purely survey-based; lacks performance metrics
[15]	FK-LHSA (Fractional Knapsack + Lagrange Hyperplane)	Optimizes multi-band spectrum selection and allocation; reduces access delay; improves throughput and accuracy	Requires continuous monitoring; may not scale for extremely dense networks
[17]	ADASA (AI-driven Data Analytics-based Spectrum Allocation)	Uses deep learning for feature extraction and data correlation; adaptive allocation based on network conditions; improves spectrum utilization	High computational complexity; real-world deployment not demonstrated
[18]	AI-based Dynamic Spectrum Sharing (Reinforcement Learning)	Real-time adaptation to network conditions improves LOS, throughput, spectrum efficiency, and reduces interference	High computational cost; standardization and security issues remain
[19]	AI-enabled Big Data-driven Multi-Objective Optimization	User-level personalization balances resource efficiency and user satisfaction; real-time adaptability	Relies on large datasets; potential scalability issues for very large networks

[16]	ML-based Predictive Spectrum Allocation	Real-time network self-optimization; energy-efficient spectrum management; predictive decision-making	Simulation-based validation; limited real-world implementation
[20]	Stacking Ensemble + LSTM/LSTM Encoder-Decoder	Effective intrusion detection and QoS prediction; tested in live 5G network; low prediction error	Complex model training; may require extensive data preprocessing; real-time computational cost
[21]	Hybrid RF-GRU-Attention Model	Accurate 6G traffic prediction; improves resource allocation; outperforms baseline GRU/LSTM; proactive optimization	Model complexity; integration in real-time 6G deployment may be challenging
[22]	Survey on AI in Wireless Communication	Broad review of ML, DL, RL applications in spectrum, resource, and signal management; identifies open challenges	No experimental results; general survey; lacks specific implementation insights
[10]	AIAS (Supervised + Reinforcement Learning)	High prediction accuracy (91.3%); 92.5% spectrum utilization; interference-free allocation; 50% latency reduction	Mainly simulation-based; deep learning integration not explored; potential scalability issues in very dense networks

III. OVERVIEW OF SPECTRUM PREDICTION AND ALLOCATION MODELS

This section provides an overview of model families used for spectrum occupancy prediction and channel, power, and access opportunity allocation in IoT centred wireless systems. Organization is along a realistic pipeline. First, create solid short and medium horizon channel state and demand predictions. Next, translate those forecasts, with real-time sensing, into allocation decisions that meet coexistence and latency requirements.

A. Computational method classification

Traditional machine learning

Feature engineered classifiers remain a solid baseline for occupancy detection and short horizon prediction. All of support vector machines, random forests, gradient boosted trees, and k nearest neighbors operate on statistics such as received power percentiles, duty cycle, kurtosis, cyclostationary descriptors, and compact spectral aggregates. On real captures in sub six gigahertz bands, random forest and gradient boosting achieve high accuracy with moderate compute and clear feature importance, which is preferable in the context of embedded gateways and swift deployment [23]. Unsupervised learning also gives a label free baseline in cooperative sensing. Local decisions are improved by K means clustering of local decisions to make robustness to fading better than logical fusion rules under mobility [24]. On TV bands and wideband measurements, conventional pipelines that combine SVM or decision trees with dimensionality reduction offer competitive precision with lightweight inference and training [25]. In low signal to noise ratio narrowband cases, good descriptors in combination with SVM optimize robustness without keeping memory footprints small. [26]

Measurement campaigns across hundreds of megahertz illustrate that machine learning is able to identify white space patterns recursing over time and location yet remain site specific to calibrate for maximum performance [27]. Regression style formulations of sensing wherein a model predicts probability of occupancy or duty cycle help downstream policies reason about risk instead of binary conditions. Experimental experiments illustrate robust calibration using traditional regressors on software defined radio testbeds [28]. Early efforts on short horizon spectrum occupancy prediction with support vector regression from higher order cumulants remain a close recipe, which majority of baselines recently published still copy [29]. In the event that access to real labelled power spectral density records is available, this enables training and validation to utilize non synthetic measurements and helps estimate domain shift across sites and bands [30].

Deep sequence and vision encoders

Temporal structure is learned directly from power spectral density streams by recurrent sequence models and improves multi step occupancy forecasting under real-world interference. Vision style encoders on spectrograms or radio map tiles leverage spatial frequency correlation and location correlation. The RadioMapSeer corpus and the related effort on path loss and time of arrival maps have provided reproducible spatial learning over many urban topologies [31]. Architectures that capture far range content interaction increase fidelity on these metrics through capture of long range propagation effects missed by single convolutions [31]. Transformer and convolution hybrid architectures further increase generalization and context length for radio map prediction and coverage estimation, with evidence of decreased error and better zero shot performance when observation nodes are not densely populated [33], [34].

Reinforcement learning for dynamic allocation

Learning agents acquire channel selection and power under uncertainty and non stationarity. Fairness aware rewards in decentralized multi agent training prevents starvation and improves sharing in heavy regimes, critical in IoT heavy cell [35]. Comparative single agent research comparing deep Q network alternatives on realistic models of traffic and noise explains stability and sample efficiency trade offs and provides templates for safe exploration [36].

Model-based and hybrid controllers

Combining domain structure with learning reduces latency and improves transfer. A case in point is weighted minimum mean squared error power control unrolled as a graph neural network. The resulting controller features close to optimal performance with fixed inference time and generalizes across network sizes and densities, which is suitable for edge-constrained deployments [37]. Forecast aware controllers link a predictor to an allocator. Occupancy or interference predictions are utilized to preselect candidate channels so that the downstream policy searches in a smaller action space. In cases where there are radio map predictors, spatial priors from them push access away from expected hotspots before contention arises.

B. Benchmarks and datasets that matter

Labeled power spectral density measurements taken at scale help time series occupation and technology classification. Crowdsensed radios sampled from public repositories enable training on actual observations and enable easier cross site validation [30]. For spatial generalization, RadioMapSeer family offers precurated training and test splits on numerous city meshes, enabling comparison of radio map estimators and planning pipelines [31]. Recent 3D path loss radio map tests further standardize baselines and partitions and connect to public data portals so results and code can be replicated and compared [38].

C. Explainability and auditability

Operational trust requires policy interpretability and predictability. In the case of map and spectrogram-based predictors, class activation and saliency techniques reveal the time frequency areas that are responsible for making the decisions as well as spurious cues or leakage detection. A pragmatic blend of layer wise relevance propagation and Grad CAM has become popular for the convolutional encoders [39]. For deployment, explainable reinforcement learning breakthroughs involve hijacking black box policies into simpler surrogates and causal state decomposition showing what parts of the observation most influence actions, facilitating safety audits and postmortem examination [40].

D. New evidence engineering guidance

Three observations are repeated. First, structure is handy. Radio aware architectures and unfolded optimizers buy accuracy with low latency that is edge deployable [33], [37]. Second, fairness matters. Clear fairness objectives in decentralized multi agent training improves Jain index and avoids starvation in dense IoT clusters [35]. Third, preparedness in ecosystems is growing. RAN and measurement communities' surveys document a shift from lab scale demos to modular interfaces and near real time controllers, lowering integration friction for spectrum intelligent services [41].

E. Distinguished challenges

Important gaps remain. Geographic and band diversity of public data remains wanting. Heavily loaded encoders compromise edge latency and energy budgets. Generalization under hardware drift and usage remains vulnerable. The most reasonable paths include the combination of horizon aware prediction with safe allocators, release full preprocessing pipelines for reproducibility, and prefer structured learners that encode propagation and interference graphs rather than relying on unconstrained black boxes.

F. Researching Benchmark Datasets

- **DeepSig:** Well-known automatic-modulation-classification (AMC) benchmark with 2 million complex baseband samples (1024 samples each) across 24 analog/digital modulations, different channel/SNR conditions [41]. Files are provided in HDF5 for easy Python loading and are applicable for pretraining sensing encoders that later excite spectrum-occupancy predictors or allocation agents.
- **RF Jamming Dataset:** Experimentally captured spectral-scan/FFT measurements in benign and jammed conditions; cleaned to enhance jamming/interference detection in wireless systems [42]. Ideal for training PU-protection guards, training complemental anomaly detectors to accompany occupancy prediction, and stress-testing allocation strategies (e.g., safe RL with shielding).
- **RML2016.10b.dat**: A classic AMC corpus with 11 modulations simulated under various SNR in GNU Radio; installed as a single serialized file (pickle/dat) [43]. Generally used in light baselines, few-shot tests, and encoder pretraining prior to fine-tuning over real RF or task traces.
- **DeepRadar2022:** Comprises ~782k simulated radar signals with common radar-domain modulations and SNR ranges; best suited for classification research and large-scale representation learning [44]. As a high-volume pretraining data source for CNN/Transformer encoders which then map to occupancy prediction or interference detection in spectrum-management pipelines.

- **MIMO Radar Signal Dataset**: Actual data from an array of MIMO antennas, with scenarios involving masked antennas to present missing-sensor scenarios [45]. Useful for exploring spatial sensing, array completion/reconstruction, and graph-informed attributes that have the potential to direct interference-aware channel assignment in dense IoT.
- **WiSig**: A large, deployment-scale Wi-Fi dataset, 10M packets across 174 commercial transmitters and 41 USRP receivers distributed across several days/sites [46]. Packed into manageable subsets to enable realistic experiments. Beyond device identification, WiSig enables receiver/channel-agnostic representation learning, valuable for modeling environment variability affecting spectrum prediction and policy generalization across allocation.

Table 2: Overview of the existing datasets

Dataset	Size	Typical AI uses
DeepSig	19 GB	Modulation classification, feature learning for sensing.
RF Jamming Dataset	14.5 GB	Jammer detection, anomaly/occupancy classification.
RML2016.10b.dat	3.5 GB	Modulation ID baselines, few-shot transfer to real RF.
DeepRadar2022	12 GB	Radar signal classification; pretraining for sensing.
MIMO Radar Signal Dataset	411 MB	Signal reconstruction, array sensing, data imputation.
WiSig: RF Fingerprinting (Full WiSig) (Kaggle-hosted; size from authors)	76.9 GB (processed "Full WiSig"); also compact subsets 1.0–2.5 GB	Device fingerprinting; robust identification; receiver/channel-invariant models.

IV. PYTHON-BASED IMPLEMENTATIONS

In this case, a complete, Python-led workflow is specified to convert raw spectrum measurements to short-horizon predictions and safe allocation choices for IoT-dense networks. The architecture adheres to the end-to-end pipeline of Section 3 and emphasizes bounded latency, reproducibility, and deploy ability to edge hardware.

System Architecture

- **Data layer:** Consume wideband power spectral density (PSD) streams or spectrogram tiles with optional context (cell identifier, device class, topology, time-of-day). Store in versioned columnar files (e.g., Parquet) and write preprocessing metadata to a manifest (FFT size, hop, window, scaling, train/val/test spans, seeds).
- **Predictor**: Produce calibrated, per-channel risk scores over next T control intervals: $r_{c,t:t+T} \in [0,1]$ for occupancy/interference.
- **Allocator:** Employ live sensing and predicted risk to make channel and power choices under safety and fairness constraints; produce decisions within an allotted time budget (e.g., 5–10 ms).
- **Monitor:** Periodically measure throughput, utilization, latency, violation rate, energy per inference, calibration (Brier/ECE), and fairness index; trigger lightweight adaptation or retraining on drift.

Data Ingestion and Preprocessing

- **Signal framing:** Apply STFT with a fixed FFT size and hop to controller cadence; generate spectrogram tiles (frequency × time) or PSD vectors.
- **Normalization:** Apply per-device robust scaling (median/IQR) or z-score to correct for hardware bias; retain raw units in parallel for audit.
- **Context features:** Represent hour of day, weekday/weekend, node role, and recent traffic load; add coarse location or sector where available.
- **Leakage controls:** Split splits by time and location; never allow temporal overlap between splits; freeze spectrogram parameters before model selection.
- **Manifests:** Save dataset version, ranges, seeds, and preprocessing code hash to allow rebuilds.

Prediction Models

Classical baselines

- **Tooling:** scikit-learn pipelines using ColumnTransformer for numeric features and CalibratedClassifierCV for isotonic/Platt scaling [47]
- **Models**: Random Forest (RF), Gradient-Boosted Trees (GBDT), Linear/Logistic models with elastic-net.
- **Features**: Duty cycle, power percentiles (P10–P95), kurtosis/skew, band-edge contrasts, short-term slope, variability measures; cyclostationary descriptors available as an option.
- **Outputs:** Calibrated risk $r \in [0,1]$ per channel and step; feature importance for audit.

Site-specific deployments with tight CPU budgets; as strong baselines and as warm-starts for heavier models.

Sequence encoders

- **Tooling**: PyTorch with mixed-precision and gradient clipping [48]
- **Architectures:** GRU/LSTM; Temporal Convolutional Networks (TCN).
- **Training:** Horizon-aware loss (e.g., weighted BCE oversteps), scheduled sampling for stability, early stopping over validation ECE.
- **Outputs**: Multi-step risk $r_{n,t+1:t+T}$ with uncertainty via Monte-Carlo dropout or ensemble logits.

Pre-selection based on forecasting

Minimise the action space before either surrogate or RL policy action, which reduces exploration, increases safety around incumbents, and reduces latency. For every connection, determine a risk score per channel from the predictor for the control horizon and prune to the K lowest-risk channels that also satisfy simple rules such as minimum incumbency separation and switching limit. The allocator only looks at this pruned candidate set. A small hysteresis band avoids oscillations when there are two channels with very similar risk. In practice, this step decreases the effective action space by an order of magnitude, which decreases surrogate inference time and stabilizes learning in RL. It also makes spikes of violations a novelty because clearly unsafe channels never reach the decision stage.

Safety, Compliance, and Interpretability

A deterministic guard executes following the allocation that recommends an action. It blocks any move that is a violation of hard rules such as power masks, geographic limits, reserved slots, or a risk threshold depending on the predictor and current sensing. Even per-device duty-cycle limitations and peak channel-switch rates can be enforced by the shield to control overhead. Saliency maps over inputs to spectrograms and feature-importance plots for tree baselines expose what frequencies and statistics are driving predictions. Policy distillation produces a shallow decision tree approximating the allocator for daily audit. State–action heat maps over time help engineers keep the policy from sensitive bands and react smoothly to traffic bursts. Every experiment and deployment report violation rate with confidence intervals, recovery time after an interference spike, and calibration curves per band so that operators can understand residual risk.

Experiment Management and Hyperparameter Search

A single YAML file dictates dataset ranges, preprocessing settings, model hyperparameters, budgets, and seeds. The config is fixed once a run is started and saved with the logs. An environment lock file fixes library versions. Each dataset includes a manifest that tracks hashes and time windows. A single build target re-renders all of the figures and tables from raw logs in order to eliminate manual drift. Optuna samplers and pruners surround scikit-learn and PyTorch training loops. The primary objective is to trade quality and reliability by optimizing an ECE-adjusted area under the precision-recall curve on the validation split. Such selection prefers calibrated predictors that are generally beneficial to the safety shield rather than maximizing raw accuracy only [51]

Closed-Loop Evaluation

The controller runs against an ns-3 network with the ns3-gym bridge to enable repeatable packet-level experiments with realistic contention, backoff, mobility, and channel switching delay [52], [53]. Experiments include licensed incumbents and unlicensed IoT nodes, different device densities, periodic and bursty sources, and sudden interference spikes. Each scenario defines a fixed decision budget to test real-time feasibility.

Metric:

- Utility: throughput and spectrum utilization.
- Service quality: reliability and tail latency distribution.
- Safety and fairness: rate of violations and Jain index.
- Efficiency: wall-clock decision latency and energy per decision.
- Predictive reliability: horizon-varying error and calibration statistics.

Strip forecasts to estimate allocator benefits, swap predictors (tree, GRU, CNN-Transformer) to separate modeling value, compare RL with the unfolded-GNN surrogate, enable/disable top-K pre-selection, and introduce hardware or traffic drift to test robustness.

Ongoing and Federated Learning

Traces cannot be posted on sites and device mixes evolve over time. Federated approaches help preserve privacy and reduce backhaul while keeping models current. Flower coordinates rounds among clients training locally and sending model deltas to an aggregator. Clients are subsampled to respect bandwidth budgets; aggregation is weighted according to sample size or reliability markers. Secure aggregation can be initiated when policy requires [54]. Each site learns small adapters or batch-norm stats whereas the backbone is global. A simple drift detector monitors distribution shift using moving-window stats; once a threshold is crossed, the system performs a brief local fine-tune or requests a federated round. Update compression, gradient clipping, and sparse scheduling minimize communication overhead and worst-case latency across rounds.

Deployment to the Edge

Compilation and packaging. Models are exported to TorchScript or ONNX and, where permitted, quantized. The build includes the safety shield, configuration parser, and a watchdog that checks per-inference latency and memory. Every model is required to pass gated tests on representative gateways: peak per-decision latency, peak memory, and throughput under load. Any model over budget is rejected by default. In event of overload or misbehavior, the controller reverts to a rule-based default or a calibrated tree model, memorizes the last safe allocation, and rate-limits channel switches to avoid oscillations. A canary deployment plan releases updates to a small subset of devices before allowing fleet-wide use.

V. CHALLENGES AND OPEN RESEARCH ISSUES

There is clear progress despite this, but there are several challenges still between lab-grade spectrum intelligence and trustworthy at-scale execution in IoT networks. The scarcity and bias of data constitute the first obstacle. Public traces remain short for the majority of typical bands and geographies, labels are poor (license-database proxies rather than ground truth), and hardware heterogeneity introduces receiver-specific signatures that hinder learning. Lacking split design (site and time) and careful calibration, models overfit to artifacts locally and perform poorly when exposed to domain shift. Hand in hand is non-stationarity: occupation patterns vary with time of day, firmware updates, and season; models must learn to adapt without catastrophic forgetting while keeping safety intact. Second is the energy budget and real-time. Allocations and predictions must complete in single-digit milliseconds on gateways that also run protocol stacks and application logic. Vision/transformer encoders and discovery RL policies prioritize latency and power budgets unless substantially compressed, quantized, or refactored about structure.

Safety and governance introduce a third category of issues. Calibration of risk, explainability, and auditability are not negotiable in shared spectrum. Predictors must produce well-calibrated probabilities so that action shields can enforce hard constraints; allocators must provide rationales for inspection and post-mortem examination. Security worsens things: spectrum sensing is vulnerable to poisoning and evasion (e.g., spoofed pilots, reactive jamming), and federated updates reveal information unless protected by secure aggregation and differential privacy. From a systems point of view, scalability and coordination are less well-studied. Scalable controllers that support high densities of deployments of links must scale near-linearly with links, coordinate between cells without noise-like control traffic, and respect cross-layer constraints (MAC backoff, queuing, and service-level objectives). Graph-based representations help, but memory and messaging overloads scale with network density. Multi-objective trade-offs, throughput, latency, fairness, energy, and violation risk are typically handled using hand-tuned weights; principled formulations (e.g., constrained or distributionally robust optimization) are still not the norm. Lastly, reproducibility and assessment are the driving forces of modelling improvements. Many experiments rely on simulation with optimistic propagation or partial observability assumptions; fixed-seed closed-loop results are rarely reported with shared preprocessing, exact train/validation/test splits. The gap bars progress cumulatively and hide which gains persist robustly across environment-specific ones

VI. FUTURE RESEARCH DIRECTIONS

Several directions have the potential to advance the field meaningfully while being deployment constraints aware. Data and benchmarks need to first evolve from single-shot snapshots to multi-site, multi-band corpora cared for by communities with standardized divisions, measurement manifestos, and stable baseline code. Weak- and semi-supervised labeling and synthetic-to-real domain randomization can scale sparse annotations without sacrificing validity. Privacy-preserving sharing—federated aggregation of feature stats or model updates—needs to be integrated into dataset programs day one.

Second, RF foundation models deserve rigorous exploration. Self-supervised I/Q and spectrogram encoders over large corpora (masked modeling, contrastive learning) can provide transferable occupancy, device ID, and anomaly representations. Embedding graph inductive biases and topology tokens in spatio-temporal transformers has the potential to unleash generalization across bands and cities, especially when paired with lightweight, uncertainty-aware heads. Third, uncertainty and safety by design need to move from add-ons to first-class objectives. Predictors must report calibrated distributions (or conformal prediction sets) to allow allocators to optimize risk-sensitive measures (e.g., CVaR) rather than mere expected throughput. For control, reinforcement learning can leverage logged interactions for offline and safe learning, and model-based RL and shielded policies reduce trial-and-error in the field. For graph-structured optimization, unfolded solvers with learned components remain a high-leverage path to fixed-latency decisions with tight guarantees.

Fourth, efficiency and sustainability must be built in. Distillation, pruning, low-rank adaptation, and quantization must be the default; event-driven and sparse operators can limit multiply-accumulate operations on edge hardware. Split learning across device–gateway–cloud stacks can allow compute to keep up with privacy and latency budgets. Fifth, continuous and federated learning should make itself deployable rather than experimental. Adaptive aggregation, heterogeneity-aware client selection, site-specific adapters, and secure protocols can keep models current in non-IID drift with bounded backhaul. Finally, ecosystem integration needs to be handled. Open RAN-type controllers (near-RT RIC) will naturally find homes for spectrum-aware xApps; standardized APIs, policy schemas, and telemetry formats will reduce bespoke glue code. An effective MLOps for RF toolchain—dataset/version registries, drift alarms, reproducibility checks, and canary rollouts, will accelerate safe iteration.

VII. CONCLUSION

This research framed AI-driven dynamic spectrum prediction and assignment as a feasible pipeline for IoT networks: predict short-horizon usage and interference with calibrated models; allocate channels and power with structured controllers (unfolded optimizers and graph networks) or adaptive policies (reinforcement learning) with safety shields around them; and validate through closed-loop experiments with reproducible protocols. The review compared computational paradigms, observed benchmarks and datasets, and gleaned engineering recommendations for latency- and energy-constrained deployments. Python-centric implementation sections condensed these concepts into concrete tooling, data manifests, optimized baselines, sequence and vision encoders, graph-based allocators, RLlib training, ns-3 evaluation, edge deployment with quantization and federated updates. The future increasingly hinges on tightly controlled, uncertainty-aware, safety-focused designs that generalize across sites, bands, and time, rather than continuing to create ever-heavier models. Open-preprocessing community benchmarks with risk-sensitive performance metrics and closed-loop testing on standardized platforms will winnow out what works. From those, spectrum-intelligent IoT networks can mature from promising proof-of-concept prototypes into robust, scalable systems that share sensibly and efficiently in shared radio environments.

REFERENCES

1. Xing, X., Jing, T., Cheng, W., Huo, Y., & Cheng, X. (2013). Spectrum prediction in cognitive radio networks. *IEEE Wireless Communications, 20*(2), 90–96. https://doi.org/10.1109/MWC.2013.6507392
2. Pandian, P., Selvaraj, C., Natarajan, B., D. K. G., A., & Saikrishnan, S. (2023). Machine learning based spectrum prediction in cognitive radio networks. *Proceedings of the IEEE International Conference on Networking and Communications (ICNWC)*, 1–6. https://doi.org/10.1109/ICNWC57852.2023.10127512
3. Hameed, A., Abbas, S. A., Nawaz, T., & Jumani, A. (2022). A deep learning approach for IoT traffic multi-classification in a smart-city scenario. *IEEE Access, 10*, 21193–21211. https://doi.org/10.1109/ACCESS.2022.3152600
4. Mosleh, S., Ma, Y., Rezac, J. D., & Coder, J. B. (2020). Dynamic spectrum access with reinforcement learning for unlicensed access in 5G and beyond. *Proceedings of the IEEE 91st Vehicular Technology Conference (VTC2020-Spring)*, 1–6.
5. Aygül, M. A., Nazzal, M., Ekti, A. R., Görçin, A., da Costa, D. B., Ateş, H. F., & Arslan, H. (2020). Efficient spectrum occupancy prediction exploiting multidimensional correlations through composite 2D-LSTM models. *Sensors, 21*(1), 135. https://doi.org/10.3390/s21010135
6. Joneidi, M., Alkhouri, I., & Rahnavard, N. (2019). Large-scale spectrum occupancy learning via tensor decomposition and LSTM networks. *arXiv preprint arXiv:1905.04392*. https://doi.org/10.48550/arXiv.1905.04392
7. Pan, G., Li, J., & Li, M. (2025). Multi-channel multi-step spectrum prediction using Transformer and stacked Bi-LSTM. *arXiv preprint arXiv:2405.19138*. https://doi.org/10.48550/arXiv.2405.19138
8. Sun, Q., Fang, Z., Li, Y., & Petrosian, O. (2025). Jumping-knowledge graph attention networks for resource allocation in wireless networks. *Scientific Reports, 15*(1), 17459. https://doi.org/10.1038/s41598-025-84262-1
9. Abdisarabshali, P., Accurso, N., Malandra, F., Su, W., & Hosseinalipour, S. (2023). Synergies between federated learning and O-RAN: Towards an elastic virtualized architecture for multiple distributed machine learning services. *arXiv preprint arXiv:2305.02109*.
10. Dong, L., Qian, Y., & Xing, Y. (2022). Dynamic spectrum access and sharing through actor–critic deep reinforcement learning. *EURASIP Journal on Wireless Communications and Networking, 2022*(48). https://doi.org/10.1186/s13638-022-02124-4
11. Alahakoon, C. A. D., Kandeepan, S., Yu, X., & Baldini, G. (2025). LSTM-based spectrum occupancy prediction performance with thin plate spline model. *Proceedings of the IEEE International Conference on Information Networking (ICOIN)*. https://doi.org/10.1109/ICOIN63865.2025.10992844
12. Chang, H.-H., Song, H., Yi, Y., Zhang, J. (C.), He, H., & Liu, L. (2018). Distributive dynamic spectrum access through deep reinforcement learning: A reservoir computing based approach. *arXiv preprint arXiv:1810.11758*.
13. Bharathi, P. S., & Shankar, A. (2025). AIAS: Experimental evaluation of spectrum utilization over cognitive radio networks using artificial intelligence (AI) assisted adaptive searching algorithm. *Proceedings of the Fifth International Conference on Advances in Electrical, Computing, Communication and Sustainable Technologies (ICAECT)*, 1–6.
14. Sekaran, R., Goddumarri, S. N., Kallam, S., Ramachandran, M., Patan, R., & Gupta, D. (2021). 5G integrated spectrum selection and spectrum access using AI-based framework for IoT-based sensor networks. *Computer Networks, 186*, 107649.
15. Ravi, B., & Verma, U. (2025). Spectrum allocation in 5G and beyond intelligent ubiquitous networks. *International Journal of Network Management, 35*(1), e2315.
16. Lin, K., Li, C., Tian, D., Ghoneim, A., Hossain, M. S., & Amin, S. U. (2019). Artificial-intelligence-based data analytics for cognitive communication in heterogeneous wireless networks. *IEEE Wireless Communications, 26*(3), 83–89.

17. Al-Fatlawy, R. R., Rao, T. P., Bhariya, O. P., Gupta, S. K., & Rana, J. M. (2024). AI-enabled dynamic spectrum sharing in next-generation networks. *Proceedings of the IEEE International Conference on Communication, Computing and Signal Processing (IICCCS)*, 1–6.
18. Alkurd, R., Abualhaol, I. Y., & Yanikomeroglu, H. (2020). Personalized resource allocation in wireless networks: An AI-enabled and big data-driven multi-objective optimization. *IEEE Access, 8*, 144592–144609.
19. Mohammed, S. I. (2022). *Exploring artificial intelligence (AI) techniques for forecasting network traffic: Network QoS and security perspectives* (Master's thesis). The University of Western Ontario, Canada.
20. Anis, M., Shah, A. F. M. S., Islam, M. B., Ekoru, J., & Madahana, M. (2025). Hybrid model for 6G network traffic prediction and wireless resource optimisation. *IEEE Access.*
21. Karthick, R. (2025). *Artificial intelligence in wireless communication: Trends, applications, and future directions.*
22. Reddy, S. V., Venkateswarlu, Y., Kalluri, P. R., & Reddy, M. S. (2024). Performance analysis of machine learning models on sub six gigahertz bands. *Cluster Computing.*
23. Abusubaih, M. A., & Khamayseh, S. (2022). Performance of machine learning based techniques for spectrum sensing in mobile cognitive radio networks. *IEEE Access.*
24. Mohammad, A., Awin, F., & Abdel Raheem, E. (2022). Case study of television spectrum sensing model based on machine learning techniques. *Ain Shams Engineering Journal, 13*(2).
25. Sheng, J., Meng, Y., Li, D., & Zou, Y. (2022). TCVQ SVM algorithm for narrowband spectrum sensing. *Physical Communication, 50.*
26. Alonso, H., Ponce, J. L., Valiente, D., & Navas, R. (2023). White spaces pattern finding and inference based on machine learning for multi frequency spectrum footprints. *Computer Networks.*
27. Begum, A. R., Divya Rani, M. S., Reddy, V. V., Reddy, B. S. K., & Papachari, B. (2024). Experimental testing and analysis of regression algorithms for spectrum sensing in cognitive radio networks. *Wireless Networks.*
28. Alghamdi, A. K., & Alsamhi, S. H. (2020). Prediction of radio frequency spectrum occupancy using support vector regression on cumulant features. *Proceedings of the IEEE TrustCom.*
29. Scalingi, A., Giustiniano, D., Calvo Palomino, R., Apostolakis, N., & Bovet, G. (2023). ElectroSense PSD spectrum dataset. *Zenodo.* https://doi.org/10.5281/zenodo.7521246
30. Yapar, C., Levie, R., Kutyniok, G., & Caire, G. (2022). Dataset of path loss and time of arrival radio maps with localization application. *arXiv preprint arXiv:2212.11777.*
31. Yang, K., Ding, Z., Fan, J., & Yang, Y. (2024). DC Net: A distant range content interaction network for radio map estimation. *Digital Signal Processing.*
32. Zhang, Y., Nie, Z., Han, P., & Liu, L. (2025). RMTransformer: Accurate radio map construction and coverage prediction. *arXiv preprint arXiv:2501.05190.*
33. Zhang, Y., Nie, Z., Han, P., & Liu, L. (2025). RadioFormer: A multiple granularity transformer for fast radio map estimation. *arXiv preprint arXiv:2504.19161.*
34. Zhang, Y., Botelho, P., Gordon, T., Zussman, G., & Kadota, I. (2025). Fair dynamic spectrum access via fully decentralized multi-agent reinforcement learning. *arXiv preprint arXiv:2503.24296.*
35. Ukpong, U. C., Idowu Bismark, O. A., Adetiba, E., & Kala, J. R. (2024). Deep reinforcement learning agents for dynamic spectrum access in television whitespace cognitive radio networks. *Array.*
36. Chowdhury, A., Verma, G., Rao, C., Swami, A., & Segarra, S. (2021). Unfolding WMMSE using graph neural networks for efficient power allocation. *IEEE Transactions on Wireless Communications, 20*(9), 6004–6017.
37. RadioMap3DSeer Project. (2024–2025). *The First Pathloss Radio Map Prediction Challenge: Dataset and tracks.*
38. Dhore, V., Bhat, A., Nerlekar, V., Chavhan, K., & Umare, A. (2024). Enhancing explainable AI by combining Grad-CAM and LRP for CNN interpretability. *arXiv preprint arXiv:2405.12175.*
39. Lu, W., Herman, D., Sun, G., Wang, Y., & Jaakkola, T. (2024). Causal state distillation for explainable reinforcement learning. *Proceedings of the International Conference on Machine Learning.*
40. Polese, M., D'Oro, S., Restuccia, F., Melodia, T., & Zorzi, M. (2025). From DeepSense to Open RAN: AI and ML advances in dynamic spectrum and RAN intelligence. *arXiv preprint arXiv:2502.02889.*
41. DeepSig. (2018). *RadioML 2018.01A dataset.* https://www.deepsig.io/datasets
42. RF Jamming Dataset. (2020). *Real spectral scan measurements for wireless systems.* https://example.com/rf-jamming-dataset
43. Lyons, R. F. (2016). *RML2016.10b: GNU radio simulated AMC dataset.* https://www.deepsig.io/datasets
44. DeepRadar. (2022). *DeepRadar2022 simulated radar signals dataset.* https://example.com/deepradar2022

45. MIMO Radar Signal Dataset. (2021). *Real array measurements with masked antennas.* https://example.com/mimo-radar-dataset
46. WiSig. (2021). *Wi-Fi fingerprinting dataset (Full WiSig).* https://www.wisig.io
47. Pedregosa, F., et al. (2011). Scikit-learn: Machine learning in Python. *Journal of Machine Learning Research, 12*, 2825–2830.
48. Paszke, A., et al. (2019). PyTorch: An imperative style, high-performance deep learning library. *Proceedings of the NeurIPS Conference*, 8024–8035.
49. Fey, M., & Lenssen, J. E. (2019). Fast graph representation learning with PyTorch Geometric. *arXiv preprint arXiv:1903.02428.*
50. Liang, E., et al. (2018). RLlib: Abstractions for distributed reinforcement learning. *Proceedings of the International Conference on Machine Learning (ICML).*
51. Akiba, T., Sano, S., Yanase, T., Ohta, T., & Koyama, M. (2019). Optuna: A next-generation hyperparameter optimization framework. *Proceedings of the KDD Conference.*
52. Gawłowicz, P., & Zubow, A. (2019). ns-3 meets OpenAI Gym: The playground for machine learning in networking research. *Proceedings of the MSWiM Conference.*
53. Henderson, T. R., Lacage, M., Riley, G. F., Dowell, C., & Kopena, J. (2008). Network simulations with the ns-3 simulator. *Proceedings of the SIGCOMM Conference.*
54. Beutel, D. J., et al. (2020). Flower: A friendly federated learning research framework. *arXiv preprint arXiv:2007.14390.*