# From Spectrum Data to Spectrum Decisions: An Agentic AI Approach for 6G Networks

Nina Rechberger
PhD Candidate Applied AI
Alma Mater Europaea University
Maribor, Slovenia
nina.rechberger@almamater.si

## Abstract

Sixth-generation networks will require spectrum management that is dynamic, context-aware, explainable, and compliant with both technical and legal requirements. This paper proposes the Agentic Spectrum Intelligence Framework (ASIF), a conceptual architecture that transforms heterogeneous spectrum data and high-level human intent into bounded, auditable, and verifiable spectrum decisions. ASIF combines specialised Monitoring, Prediction, Decision, Compliance, and Explanation Agents coordinated by a large-language-model-based Spectrum Assistant. The Spectrum Assistant interprets intent, supports orchestration, and produces human-readable explanations, but does not directly allocate spectrum resources. A distinguishing feature of the framework is the Compliance Agent, which evaluates proposed actions against licensing conditions, coexistence rules, harmful-interference limits, service priorities, and governance obligations before execution. The paper further introduces a regulatory and data-governance layer for translating authoritative requirements into machine-processable constraints, validation rules, and audit records. An urban emergency-connectivity use case illustrates the compliance-controlled closed-loop workflow. Finally, the paper proposes a future validation plan covering technical performance, decision quality, governance, explainability, auditability, and human oversight. ASIF is presented not as a validated operational system, but as a research framework for investigating accountable, policy-governed, and human-supervised agentic spectrum management in future 6G networks.

## Keywords

Radio Spectrum Management, 6G Networks, Agentic AI, Explainable AI, Spectrum Governance

Permission to make digital or hard copies of part or all of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for third-party components of this work must be honored. For all other uses, contact the owner/author(s).
*Information Society 2026, 5–9 October 2026, Ljubljana, Slovenia*
© 2025 Copyright held by the owner/author(s).
http://doi.org/10.70314/is.2026.skui.54

## 1 Introduction

Radio spectrum is the range of electromagnetic frequencies used by wireless communication systems. Because several services may seek access to the same limited frequency resources, spectrum management determines who may use a frequency band, where, when, at what power level, and under which interference and licensing constraints. In this paper, a spectrum decision means a bounded operational choice, such as approving temporary spectrum sharing, reconfiguring a network slice, changing a service priority, or taking no action. Sixth-generation, or 6G, networks refer here to the emerging mobile systems envisioned for the IMT-2030 era, in which AI is expected to become an integral part of network control and optimisation [1].

Existing AI-based approaches can predict spectrum occupancy, detect interference, and optimise resource allocation [2-4]. However, a technically optimal action is not necessarily legally permissible or operationally acceptable. It may violate a licence condition, exceed an interference threshold, disregard the priority of a public-interest service, or lack sufficient evidence for audit. The research problem is therefore not only how to optimise spectrum use, but how to ensure that an AI-supported spectrum decision is technically feasible, compliant with authoritative constraints, explainable to a human decision-maker, and auditable after execution. This paper addresses the following research question: How can heterogeneous spectrum data and high-level human intent be transformed into technically feasible, standards-constrained, explainable, and auditable spectrum actions without delegating unrestricted decision authority to an LLM? The objective is to define a conceptual and testable agentic AI architecture for compliance-controlled spectrum decision-making. The expected result is not a validated allocation algorithm, but a precise specification of the actors, information flows, decision gates, outputs, and validation criteria required for such a system.

The paper makes three contributions. First, it defines the transformation from spectrum data and human intent to bounded spectrum actions. Second, it introduces a dedicated Compliance Agent as an ex ante control gate between technical recommendation and execution. Third, it demonstrates the workflow through an urban emergency-connectivity scenario and derives a comparative validation plan.

## 2 Related Work

AI-enabled radio-frequency spectrum management has advanced spectrum sensing, prediction, interference management, resource allocation, and dynamic sharing [2, 3, 4]. Research on O-RAN and AI-native networks increasingly embeds intelligence into network control and orchestration [5]. LLM-based and intent-based approaches translate high-level objectives into operational policies and coordinate network functions [6, 7, 8], while recent surveys and architectural proposals extend these capabilities towards tool-using and multi-agent systems for 5G and 6G operations [9, 10].

ASIF does not claim to be the first application of agentic AI to 5G or 6G network management. Its narrower contribution is the explicit separation of probabilistic intent interpretation from deterministic compliance verification and controlled execution in spectrum decision-making. Unlike architectures in which an LLM or coordinating agent may directly select or trigger network actions, ASIF places a dedicated Compliance Agent and, where required, a human approval gate between technical recommendation and execution. The claimed novelty therefore lies in the decision-control structure and its spectrum-specific governance model, not in the general use of LLMs or multiple agents.

## 3 Problem Definition and Decision Model

### 3.1 Intended Users and Decision Authority

ASIF is intended for use by a mobile network operator, a shared-spectrum operator, or another authorised spectrum user. A spectrum regulator is not necessarily the day-to-day system operator, but may define authoritative rules, audit requirements, reporting obligations, and categories of actions that require prior approval. Network engineers provide operational objectives and supervise technical performance, while compliance specialists maintain validated rule sets and review escalated cases. ASIF therefore supports, rather than replaces, institutional decision-makers.

Decision authority is intentionally distributed. The Spectrum Assistant interprets intent and coordinates tasks but cannot approve or execute a spectrum action. The Monitoring and Prediction Agents provide evidence and forecasts. The Decision Agent generates and ranks technically feasible alternatives. The Compliance Agent approves, modifies, blocks, or escalates an alternative according to validated constraints. The execution component may implement only an approved action and only within its authorised scope.

### 3.2 From Human Intent to a Spectrum Action

High-level human intent denotes a desired operational outcome expressed without specifying the underlying network configuration. For example, an operator may request, “Maintain ultra-reliable emergency connectivity in the city centre for the next 30 minutes,” rather than manually selecting a frequency band, power level, network slice, and interference threshold. ASIF decomposes such an intent into measurable objectives, scope, constraints, and unresolved parameters that specialised agents can evaluate.

The workflow has five stages. Spectrum Data integrates occupancy, quality of service, interference, geolocation, licensing conditions, and network-slice status. Knowledge Generation applies prediction and anomaly detection to identify congestion risks and sharing opportunities [2-4]. Intent Understanding converts the desired outcome into a structured objective [6-8]. Autonomous Decision-Making generates and ranks alternatives, including temporary sharing, slice reconfiguration, priority adjustment, and no action. Controlled Execution implements only an approved, fully parameterised action and returns evidence to the monitoring loop.

## 4 ASIF Architecture

ASIF contains five specialised agents, an LLM-based Spectrum Assistant, and a separate deterministic execution component. Table 1 specifies each component's input, processing function, output, and execution authority.

**Table 1: ASIF Components, Information Flows, Outputs, and Execution Authority**

| Component | Input | Processing | Output | May execute ? |
|---|---|---|---|---|
| Spectrum Assistant | Human intent, context | Intent decomposition and orchestration | Structured objective and task plan | No |
| Monitoring Agent | Telemetry, occupancy, QoS, geolocation | State estimation and anomaly detection | Current spectrum-state representation | No |
| Prediction Agent | Current and historical state | Demand and interference forecasting | Forecast with uncertainty | No |
| Decision Agent | State, forecast, objectives | Candidate generation and ranking | Ranked feasible alternatives | No |
| Compliance Agent | Candidate action, authoritative rules, licence data | Deterministic and policy-based validation | Approve, modify, block, or escalate | No |
| Explanation Agent | Evidence, alternatives, rules, outcome | Trace synthesis | Human- and machine-readable record | No |
| Execution component | Approved bounded action | Network actuation | Execution result and feedback | Yes |

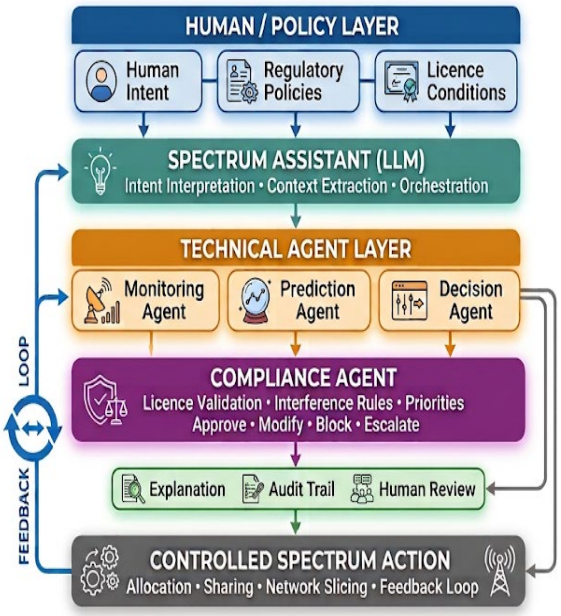


**Figure 1. ASIF workflow from spectrum data and human intent to an approved spectrum action. The Decision Agent proposes candidate actions; the Compliance Agent approves, modifies, blocks, or escalates them. Only an approved and fully parameterised action may be passed to the network execution component.**

# 5 Reliability and Compliance Controls

## 5.1 Reliability Boundaries for the LLM

The probabilistic character of an LLM is incompatible with unrestricted authority over safety-, licence-, or interference-critical network actions. ASIF therefore treats the LLM as a non-authoritative interface and orchestration component. Its outputs are proposals that must be converted into a predefined structured schema and validated before entering the decision pipeline.

Reliability is enforced through five boundaries. First, the LLM retrieves regulatory and technical information only from approved, versioned sources. Second, its output must conform to a typed intent schema containing the objective, geographical scope, temporal scope, service class, confidence, and unresolved ambiguities. Third, deterministic validators reject missing, inconsistent, or out-of-range parameters. Fourth, the Compliance Agent evaluates candidate actions against machine-readable licence conditions, interference thresholds, coexistence rules, and approval requirements. Fifth, cases involving insufficient evidence, conflicting rules, low confidence, or material risk are escalated for human approval.

Consequently, the LLM may explain, decompose, retrieve, and coordinate, but it may not create binding rules, resolve legal ambiguity autonomously, or directly execute a spectrum action. Standards adherence is handled by validated constraints and deterministic interfaces outside the generative model.

## 5.2 Regulatory and Data-Governance Layer

Autonomous spectrum decisions require legally meaningful evidence, not only technically useful telemetry. The European Electronic Communications Code frames spectrum management around effective and efficient use under objective, transparent, non-discriminatory, and proportionate criteria [11]. The EU AI Act establishes requirements relevant to risk management, data governance, transparency, logging, and human oversight [12]. The Data Act addresses fair access to and use of data, interoperability, and portability [13]. The proposed Digital Networks Act further stresses shared spectrum use and dynamic information on usage opportunities [14].

These requirements must be translated into versioned constraints, validation rules, controlled vocabularies, knowledge objects, and audit fields. For every proposed action, the Compliance Agent identifies the jurisdiction, frequency band, rights holder, service class, geographical and temporal scope, permitted power or interference limits, and required approval level. It then approves, modifies, blocks, or escalates the action while preserving the evidence, applicable rule version, unresolved ambiguity, and justification required for human accountability.

# 6 Urban Emergency-Connectivity Use Case

Consider an authorised mobile network operator that must maintain emergency-service connectivity in a congested urban area during a defined event. The operator submits the intent: “Maintain the required emergency-service quality of service within the specified area and time window while avoiding harmful interference to incumbent users.” The Spectrum Assistant converts the request into a structured objective containing the service class, geographical boundary, validity period, minimum service requirements, and unresolved parameters. The Monitoring Agent produces the current spectrum-state representation, and the Prediction Agent estimates near-term demand and interference risk. The Decision Agent generates four alternatives: temporary spectrum sharing, network-slice reconfiguration, service-priority adjustment, and no action. For each alternative, it returns the expected technical effect and uncertainty rather than a final command. The Compliance Agent verifies the applicable rights of use, authorised frequency range, geographical and temporal scope, power and interference limits, coexistence obligations, service-priority rules, and approval requirements [11, 14]. An alternative with missing evidence or incompatible constraints is blocked or escalated. Only an approved, fully parameterised action is passed to the execution component. The resulting state is continuously monitored, and adjustment or rollback is initiated if an authorised threshold is exceeded. The Explanation Agent records the evidence, evaluated alternatives, applicable rule versions, outcome, and justification.

## 7 Evaluation Plan and Research Challenges

ASIF remains a conceptual framework. Future work should implement a controlled proof of concept based on traceable spectrum data, a restricted rule set, and reproducible spectrum-sharing or network-slicing scenarios. Evaluation should cover: (i) technical performance, including utilisation, interference, latency, throughput, and service quality; (ii) decision quality, including feasibility, response time, stability, and adaptability; (iii) governance, including rule coverage, false approvals and rejections, trace completeness, explainability, and auditability; and (iv) human oversight, including reviewer agreement, correction effort, and trust calibration.

The comparison should include conventional rule-based management, AI-assisted optimisation without a Compliance Agent, and the full ASIF workflow. Ablation tests should remove the Compliance Agent, authoritative retrieval, and the human gate separately. Stress tests should vary data quality, conflicting rules, distribution shift, adversarial inputs, and interference conditions. Results should report uncertainty and failure cases, not only average performance.

Three implementation challenges remain. First, ASIF requires timely, interoperable, representative, and secure spectrum data. Second, machine-readable constraints must preserve provenance, versioning, exceptions, and escalation procedures. Third, autonomy must be proportional to risk: low-risk optimisation may be automated, while actions affecting interference protection, competition, public-interest services, or rights of use require prior approval. Security, model error, adversarial inputs, and failure propagation between agents must also be treated as system-level risks [11, 12].

## 8 Conclusion

ASIF connects spectrum evidence, intent interpretation, multi-agent reasoning, compliance verification, explanation, and controlled execution. Its main contribution is the separation of probabilistic LLM-based assistance from authoritative compliance validation and bounded network execution. The framework does not replace operator or regulatory responsibility; it defines how technical recommendations can be checked, explained, approved, and audited before implementation. The proposed architecture, concrete urban workflow, and validation plan provide a focused basis for empirical research on accountable agentic spectrum management in future 6G networks.


## Acknowledgments

This work was encouraged by the University of Alma Mater Europaea and the Doctoral Programme in Applied Artificial Intelligence.


## References


[1] International Telecommunication Union Radiocommunication Sector (ITU-R). 2023. Recommendation ITU-R M.2160: Framework and Overall Objectives of the Future Development of IMT for 2030 and Beyond. International Telecommunication Union, Geneva, Switzerland.

[2] Bushra Sabir, Shuiqiao Yang, David Nguyen, Nan Wu, Alsharif Abuadbba, Hajime Suzuki, Shangqi Lai, Wei Ni, Ding Ming, and Surya Nepal. 2024. *Systematic Literature Review of AI-enabled Spectrum Management in 6G and Future Networks*. CoRR abs/2407.10981 (2024). DOI: 10.48550/arXiv.2407.10981

[3] Xuelin Cao, Bo Yang, Kaining Wang, Xinghua Li, Zhiwen Yu, Chau Yuen, Yan Zhang, and Zhu Han. 2024. *AI-Empowered Multiple Access for 6G: A Survey of Spectrum Sensing, Protocol Designs, and Optimizations*. CoRR abs/2406.13335 (2024). DOI: 10.48550/arXiv.2406.13335.

[4] Ying-Chang Liang. 2020. Dynamic Spectrum Management: From Cognitive Radio to Blockchain and Artificial Intelligence. Springer, Singapore.

[5] O-RAN Alliance. 2024. O-RAN Native AI Architecture Description. O-RAN Alliance Research Report. https://www.o-ran.org/research-reports/o-ran-native-ai-architecture-description

[6] Abdulhalim Dandoush, Viswanath Kumarskandpriya, Mueen Uddin, and Usman Khalil. 2024. *Large Language Models Meet Network Slicing Management and Orchestration*. CoRR abs/2403.13721 (2024). DOI: 10.48550/arXiv.2403.13721

[7] Abdelkader Mekrache, Adlen Ksentini, and Christos Verikoukis. 2024. *Intent-Based Management of Next-Generation Networks: an LLM-centric Approach*. IEEE Network 38, 5 (2024), 29-36. DOI: 10.1109/MNET.2024.3420120

[8] Lam Dinh, Sihem Cherrared, Xiaofeng Huang, and Fabrice Guillemin. 2025. *Towards End-to-End Network Intent Management with Large Language Models*. arXiv:2504.13589.

[9] Mohamed Amine Ferrag, Abderrahmane Lakas and Merouane Debbah. 2026. *6G Needs Agents: Toward Agentic AI-Native Networks for Autonomous Intelligence*. CoRR abs/2605.01546. https://doi.org/10.48550/arXiv.2605.01546.

[10] Mazene Ameur, Abdelkader Mekrache, Bouziane Brik, and Adlen Ksentini. 2026. *LLM-Powered Agentic AI for 5G/6G Networks: A Tutorial and Survey on Architectures, Protocols, and Standardization*. CoRR abs/2607.16066 (2026). DOI: 10.48550/arXiv.2607.16066

[11] European Parliament and Council of the European Union. 2018. Directive (EU) 2018/1972 of 11 December 2018 Establishing the European Electronic Communications Code. Official Journal of the European Union.

[12] European Parliament and Council of the European Union. 2024. Regulation (EU) 2024/1689 of 13 June 2024 Laying Down Harmonised Rules on Artificial Intelligence (Artificial Intelligence Act). Official Journal of the European Union.

[13] European Parliament and Council of the European Union. 2023. Regulation (EU) 2023/2854 of 13 December 2023 on Harmonised Rules on Fair Access to and Use of Data (Data Act). Official Journal of the European Union.

[14] European Commission. 2026. Proposal for a Regulation for the Digital Networks Act (DNA). European Commission, Brussels.