
Reasoning Meets Representation: Envisioning Neuro-Symbolic Wireless Foundation Models

Jaron Fontaine

Department of Information Technology
Ghent University, Belgium
jaron.fontaine@ugent.be

Mohammad Cheraghinia

Department of Information Technology
Ghent University, Belgium
mohammad.cheraghinia@ugent.be

John Strassner

Americas Standards and Industry Division
Futurewei Technologies, USA
john.sc.strassner@futurewei.com

Adnan Shahid

Department of Information Technology
Ghent University, Belgium
adnan.shahid@ugent.be

Eli De Poorter

Department of Information Technology
Ghent University, Belgium
eli.depoorter@ugent.be

Abstract

Recent advances in Wireless Physical Layer Foundation Models (WPFMs) promise a new paradigm of universal Radio Frequency (RF) representations. However, these models inherit critical limitations found in deep learning such as the lack of explainability, robustness, adaptability, and verifiable compliance with physical and regulatory constraints. In addition, the vision for an AI-native 6G network demands a level of intelligence that is deeply embedded into the systems and is trustworthy. In this vision paper, we argue that the neuro-symbolic paradigm, which integrates data-driven neural networks with rule- and logic-based symbolic reasoning, is essential for bridging this gap. We envision a novel Neuro-Symbolic framework that integrates universal RF embeddings with symbolic knowledge graphs and differentiable logic layers. This hybrid approach enables models to learn from large datasets while reasoning over explicit domain knowledge, enabling trustworthy, generalizable, and efficient wireless AI that can meet the demands of future networks.

1 Introduction

Machine Learning (ML) for wireless systems has largely evolved along two paths: symbolic approaches grounded in domain knowledge, and sub-symbolic (data-driven, or neural) methods fueled by large-scale training. Symbolic methods offer transparency and explainability, which provides compliance with physical or regulatory constraints, yet struggle with addressing high levels of complexity and adaptability. Neural methods excel at learning complex patterns from raw signal data but typically function as "black boxes," limiting their reliability and trustworthiness.

Recently, large foundation models pre-trained on diverse Radio Frequency (RF) data are proposed to create universal representations [Fontaine et al., 2024, Guo et al., 2025], which is a significant step forward. Industry initiatives are already aiming for Wireless Physical Layer Foundation Models

(WPFMs) to achieve "semantic RF understanding" Fontaine and Shahid [2025] and can transfer knowledge across tasks such as modulation recognition, interference classification, sensing, resource management, and localization. However, such neural network foundation models still struggle with addressing trust and explainability.

We argue that a neuro-symbolic approach, which integrates data-driven neural networks with rule- and logic-based symbolic reasoning, is crucial to build the next generation of wireless foundation models Wan et al. [2024]. In this paper, we propose the principles of a Neuro-symbolic (NeSy) AI framework, which combines the perceptual power of deep learning with formal reasoning, to the unique constraints and opportunities of the wireless physical layer. By integrating the flexibility of neural networks with the explainability and formal guarantees of symbolic systems, we can create foundation models that are not only powerful but also trustworthy. The key motivations and the distinct advantages of our proposed solution over existing approaches are summarized in Table 1.

Metric	Symbolic Methods (Bhuyan et al., 2024)	Neural Methods (Schwalbe and Finzel, 2024)	Neuro-symbolic methods with WPFM interaction (Our vision)
<i>Explainability</i>	High (explicit rules & formulae)	Low ("black box")	High (symbolic reasoning)
<i>Scalability</i>	Low (brittle; rule explosion, complex logic)	High (learns from data)	High (leverages neural scalability)
<i>Adaptability</i>	Low (requires manual editing)	High (via ICL, etc.)	Very High (reasons over knowledge + few-shot)
<i>Constraint Compliance</i>	High (hard-coded)	Low (unable to prove)	High (verifiable via symbolic proofs)
<i>Data Efficiency</i>	High (uses domain knowledge)	Low (needs massive datasets)	Medium (priors reduce data needs)

Table 1: Comparison of Symbolic, Neural, and Neuro-symbolic (WPFM-integrated) methods across key metrics.

2 A Neuro-Symbolic Framework for Wireless AI

In this section, we discuss the proposed framework for wireless problems utilizing WPFM and neuro-symbolic methods which is illustrated in Figure 1.

2.1 WPFM: The Neural Perception Engine

The field of wireless communications is struggling with an unsustainable paradigm of training specialized models for each specific task. Focusing on one task at a time is expensive and slow, from gathering large datasets to repetitive deployment. Recent advances in RF foundation models demonstrate the feasibility of universal representations trained on diverse datasets [Cheraghinia et al., 2025, Zhou et al., 2025].

The pre-training of such foundation models, typically Transformers, does not rely on extensively labeled datasets but can learn from unlabeled data in a self-supervised way. To do so, masking or next-sample prediction can be applied to learn rich physical-layer embeddings. The input data of the WPFM can be represented as raw time series, e.g. in-phase quadrature (IQ) samples, channel impulse response (CIR), channel state information (CSI), while tokenization can be performed using patching. To retain the position of the patches during the self-attention mechanism, positional or learned encoding schemes are applied. After the WPFM is finetuned on a large variety of RF data, a lightweight finetuning process helps the model to quickly adapt to new downstream tasks (classification or predictions) with minimal labeling requirements.

In this paper, we propose that these foundation models can be further improved by producing embeddings (sub-symbolics), instead of traditional label classifications or predictions, anticipating the need for interpretable decision-making and reusable information for different tasks from a single

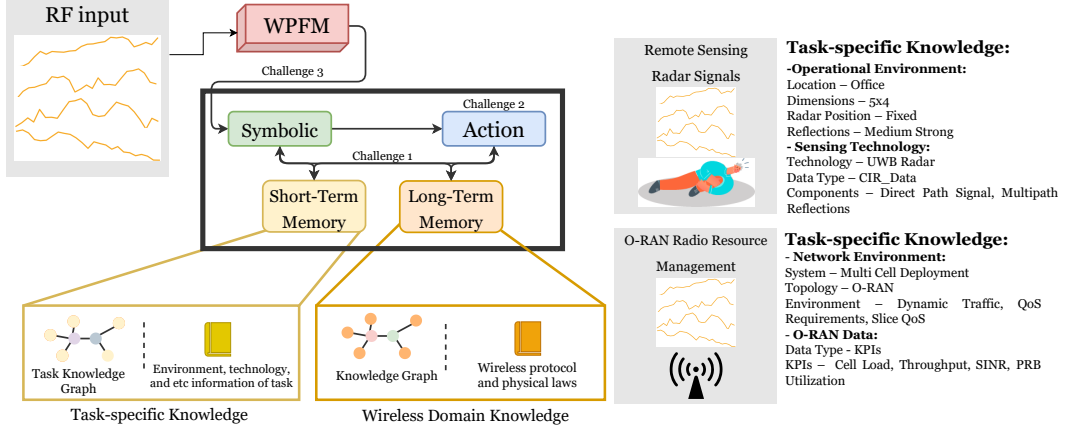


Figure 1: A Neuro-Symbolic Framework for Wireless AI: that comprises of a WPFM that embeds RF signals to be used by the symbolic component. The symbolic component utilizes memories, which include task-specific and general information from the wireless domain, to generate interpretable symbols. Using symbolic reasoning, actions can be performed, e.g., fall detection or O-RAN radio resource management, which are included as examples. In the right part of this figure, we included task-specific knowledge examples on top of general wireless domain knowledge, which aim to improve the accuracy and explainability of the system simultaneously.

model. These foundation models act as the "eye" of the process, receiving RF modalities and producing input to the symbolic module.

2.2 A Modular Neuro-Symbolic Wireless AI Architecture

To overcome the limitations of a purely neural approach, we propose extending the WPFM within a modular neuro-symbolic architecture, adopting a neuro-symbolic pipeline Wan et al. [2024], Liang et al. [2025]. This architecture cleanly separates the tasks of perception and reasoning. The neural perception transforms raw RF signals into a dense embedding. Next, the symbolic module integrates these embeddings with task-specific and wireless domain knowledge. This modular design grounds the neural network's pattern recognition in verifiable logic, enabling explainability and constraint enforcement.

The (neuro) symbolic reasoning pipeline comprises of three components: (1) ontologies to provide a shared semantic schema, (2) a knowledge graph, which is a concrete instance of the ontologies, to represent a structured world model, and (3) differentiable logic to make the symbolic reasoning process itself trainable via gradient-based methods. By representing logical rules as continuous functions, it becomes possible to backpropagate gradients from the symbolic layer through to the neural WPFM, enabling end-to-end training of the entire hybrid system. This transformation is often accomplished using techniques from fuzzy logic, where discrete operators like AND and OR are replaced by continuous, differentiable functions known as t-norms Van Krieken et al. [2022].

3 Research Challenges & Open Questions

The integration of symbolic reasoning with large-scale neural embeddings for wireless systems presents several unique research challenges.

Challenge 1: Knowledge Representation for the Physical Layer. A fundamental challenge is in creating a formal, machine-readable knowledge base for the wireless domain. It is not enough to just digitize standards documents; it requires encoding the continuous laws of physics (e.g., Maxwell's equations, channel fading models) and the discrete logic of communication protocols (e.g., MAC layer state transitions) into a unified formalism. This challenge needs to answer if a wireless knowledge graph can represent both continuous and discrete logics in a way that is compliant with robust symbolic reasoning.

Challenge 2: Real-Time Neuro-Symbolic Inference. Wireless physical layer systems operate under strict, microsecond-level constraints. While WPFMs are computationally intensive, their inference latency can be improved using model compression techniques such as knowledge distillation to train smaller, faster models, quantization to reduce numerical precision for hardware acceleration, and pruning to eliminate redundant parameters. However, symbolic reasoning can also be slow. As such, it becomes a major challenge to achieve real-time inference with the entire hybrid model. This leads to the question whether we can leverage techniques like differentiable logic circuits, which can be compiled directly to hardware (FPGAs/ASICs), to create a reasoning layer that meets the demanding deadlines of physical layer control loops.

Challenge 3: Bridging the Sub-symbolic (embeddings) to Symbolic Gap. A core challenge in all neuro-symbolic AI is the inherent mismatch in knowledge representation. The self-attention mechanisms of foundation models create distributed, non-local representations (sub-symbolic) that have complex interactions. In contrast, symbolic AI represents knowledge explicitly using discrete, localized structures like symbols and rules. To overcome this challenge, the gap between the continuous, distributed representations of neural networks and the discrete, structured representations of symbolic logic needs to be bridged. Finding semantically meaningful and computationally effective ways to translate information across this gap is a research problem motivating this work.

This vision opens several critical questions to be addressed the research community. First, what is the inherent trade-off between the expressive richness of the symbolic representation and the real-time inference capability of a NeSy-WPFM? Does a Pareto frontier exist between reasoning depth and latency in this domain? Second, beyond task-specific accuracy, how do we evaluate the "trustworthiness" of a NeSy-WPFM? What novel benchmarks and methodologies are required to quantify improvements in interpretability and robustness to out-of-distribution scenarios, such as novel interference types or zero-day exploits at the physical layer?

4 Impact on Applications

The development of neuro-symbolic WPFMs will be a foundational step toward realizing the vision of AI-native 6G Yang et al. [2024]. This vision is about designing networks where AI is "natively integrated from the ground up" into the core fabric of the system. A NeSy-WPFM solution represents a concrete instantiation of this principle at the physical layer, providing the Radio Access Network with the capability for reasoning, self-explanation, and verifiable compliance.

This framework lays the foundation for applications such as:

Trustworthy and Verifiable Wireless AI: Enabling systems that can explain their decisions (e.g., "I am allocating this resource block because of predicted interference from source X, which violates regulatory mask Y"), generalize to unseen network conditions, and guarantee compliance to physical laws and spectrum policies.

Cognitive and Autonomous 6G RAN: Embedding reasoning directly into physical-layer intelligence, moving from today's reactive optimization to proactive, self-evolving networks that can autonomously manage resources, mitigate interference, and diagnose faults with minimal human intervention.

Unified Sensing and Communication: Integrating sensing and communication under a single, reasoning-aware solution unlocks new cross-domain applications, such as remote sensing in health-care, localization for radio resource management.

5 Conclusion and future work

By combining universal neural RF representations with symbolic reasoning, we can overcome the limitations of purely data-driven approaches and move toward trustworthy, generalizable, and efficient systems. This hybrid approach is essential for realizing the vision of an AI-native wireless system where intelligence is robustly embedded into the system.

A primary direction for future work is to investigate these neuro-symbolic capabilities to empower agentic AI. The reasoning offered by this framework can serve as the cognitive engine for autonomous agents capable of network management, self-diagnosis, and optimization with minimal human intervention. However, a requirement for such advanced agents is a standardized, machine-readable

knowledge base. Therefore, another future direction can be developing a standardized, open-source wireless knowledge graph. This involves encoding the continuous laws of physics and the discrete logic of communication protocols into a unified structure, providing the necessary information for robust symbolic reasoning.

References

- Bikram Pratim Bhuyan, Amar Ramdane-Cherif, Ravi Tomar, and TP Singh. Neuro-symbolic artificial intelligence: a survey. *Neural Computing and Applications*, 36(21):12809–12844, 2024.
- Mohammad Cheraghinia, Eli De Poorter, Jaron Fontaine, Merouane Debbah, and Adnan Shahid. Foundation Model for Wireless Technology Recognition Using IQ Timeseries, 2025. URL <https://arxiv.org/abs/2505.19390>.
- Jaron Fontaine and Adnan Shahid. Reimagining Wireless Intelligence: A Foundation Model for the Physical Layer. Technical report, Alliance for Telecommunications Industry Solutions (ATIS), August 2025. URL <https://atis.org/reimagining-wireless-intelligence-a-foundation-model-for-the-physical-layer/>.
- Jaron Fontaine, Adnan Shahid, and Eli De Poorter. Towards a Wireless Physical-Layer Foundation Model: Challenges and Strategies. In *2024 IEEE International Conference on Communications Workshops (ICC Workshops)*, pages 1–7, 2024. doi: 10.1109/ICCWorkshops59551.2024.10615509.
- Jiajia Guo, Yiming Cui, Shi Jin, and Jun Zhang. Large AI Models for Wireless Physical Layer, 2025. URL <https://arxiv.org/abs/2508.02314>.
- Baoyu Liang, Yuchen Wang, and Chao Tong. AI Reasoning in Deep Learning Era: From Symbolic AI to Neural-Symbolic AI. *Mathematics*, 13(11), 2025. ISSN 2227-7390. doi: 10.3390/math13111707. URL <https://www.mdpi.com/2227-7390/13/11/1707>.
- Gesina Schwalbe and Bettina Finzel. A comprehensive taxonomy for explainable artificial intelligence: a systematic survey of surveys on methods and concepts. *Data Mining and Knowledge Discovery*, 38(5):3043–3101, 2024.
- Emile Van Krieken, Erman Acar, and Frank Van Harmelen. Analyzing differentiable fuzzy logic operators. *Artificial Intelligence*, 302:103602, 2022.
- Zishen Wan, Che-Kai Liu, Hanchen Yang, Ritik Raj, Chaojian Li, Haoran You, Yonggan Fu, Cheng Wan, Sixu Li, Youbin Kim, Ananda Samajdar, Yingyan Lin, Mohamed Ibrahim, Jan M. Rabaey, Tushar Krishna, and Arijit Raychowdhury. Towards Efficient Neuro-Symbolic AI: From Workload Characterization to Hardware Architecture. *IEEE Transactions on Circuits and Systems for Artificial Intelligence*, 1(1):53–68, 2024. doi: 10.1109/TCASAI.2024.3462692.
- Yang Yang, Jianjun Wu, Tianjiao Chen, Chenghui Peng, Jun Wang, Juan Deng, Xiaofeng Tao, Guangyi Liu, Wenjing Li, Li Yang, Yufeng He, Tingting Yang, A. Hamid Aghvami, Frank Eliassen, Schahram Dustdar, Dusit Niyato, Wanfei Sun, Yang Xu, Yannan Yuan, Jiang Xie, Rongpeng Li, and Cuiqin Dai. Task-Oriented 6G Native-AI Network Architecture. *IEEE Network*, 38(1): 219–227, 2024. doi: 10.1109/MNET.2023.3321464.
- Fuhui Zhou, Chunyu Liu, Hao Zhang, Wei Wu, Qihui Wu, Derrick Wing Kwan Ng, Tony Q. S. Quek, and Chan-Byoung Chae. Spectrumfm: A foundation model for intelligent spectrum management, 2025. URL <https://arxiv.org/abs/2505.06256>.