RA-NN: A NEURAL SYSTEMS FORMULATION FOR AI-NATIVE RADIO ACCESS NETWORKS

Abstract

Radio access networks are tightly coupled physical systems whose behavior emerges from radio propagation, interference, mobility, and traffic dynamics. Observability is indirect and mediated through user equipment, resulting in partial, distributed, and time-varying evidence of network state. Existing abstractions for RAN control and optimization assume fixed structure and task-specific decision logic, limiting their suitability as a foundation for AI-native operation.

This whitepaper defines the Radio Access Neural Network (RA-NN), a neural systems formulation in which the radio access network itself is treated as a learnable system. In the RA-NN, coverage and inter-cell influence are learned directly from UE observations and maintained as intrinsic system state rather than predefined topology. User experience is represented distributionally over per-UE link states, enabling coherent network-level reasoning, coordination, and policy learning under non-stationarity. By aligning learning with the physical interaction structure of radio networks, the RA-NN establishes a concrete architectural foundation for AI-native RAN operation.

Michael Chiaramonte

1/5/2026

1. REQUIREMENTS FOR AI-NATIVE RADIO NETWORKS

Radio access networks are tightly coupled physical systems whose behavior emerges from radio propagation, interference, spatial reuse, mobility, and traffic dynamics. These effects are inseparable at the system level and cannot be decomposed into independent or static control problems.

Network observability is indirect. Observable evidence of network behavior is mediated through user equipment, which provides localized, partial, and time-varying measurements reflecting the superposition of signals from multiple cells.

To operate coherently under these conditions, observations must be integrated into an internal representation that captures relevance, influence, and interaction across the network. Without such a representation, intelligence remains fragmented and incapable of coherent coordination.

Applying intelligence at the level of individual control functions fragments decision-making across task-specific abstractions. Actions taken by one function alter radio conditions for others, yet task-level intelligence lacks a mechanism to reason holistically about these interactions. Without a common substrate that represents how influence propagates across the network, any intelligent algorithms cannot generalize across functions or scale to network-level coordination.

2. THE RADIO ACCESS NEURAL NETWORK (RA-NN) FORMULATION

Transforming the radio access network into an AI-native system requires a formulation in which learning is aligned with the structure through which network behavior arises.

The Radio Access Neural Network (RA-NN) defines this formulation. The RA-NN is a foundational system architecture that embeds learning directly into the representation of network structure, interaction, and influence. It is not defined by a particular algorithm, training procedure, or deployment topology. Instead, it specifies how the radio access network is represented such that learning, coordination, and control can operate coherently across the system. In this formulation, the network is not treated as an external system to be assisted by intelligence, but as the object of learning itself.

2.1 THE RADIO ACCESS NETWORK AS A PHYSICAL INTERACTION GRAPH

A radio access network is a physical interaction system whose behavior is governed by radio propagation, interference, spatial reuse, and mobility. Network behavior emerges from the superposition of these effects and cannot be decomposed into independent or static components. Changes to transmission parameters, antenna geometry, or user distribution modify the surrounding radio environment and influence other elements of the network. These influences are continuous, asymmetric, and governed by universal physical factors such as path loss, antenna patterns, and interference coupling.

Accordingly, the RAN is represented in the RA-NN as an interaction graph in which edges encode influence rather than explicit configuration or control relationships. This interaction structure is sparse due to physical locality and evolves as users and environments change. These properties arise from the physics of wireless communication and are treated as architectural constraints.

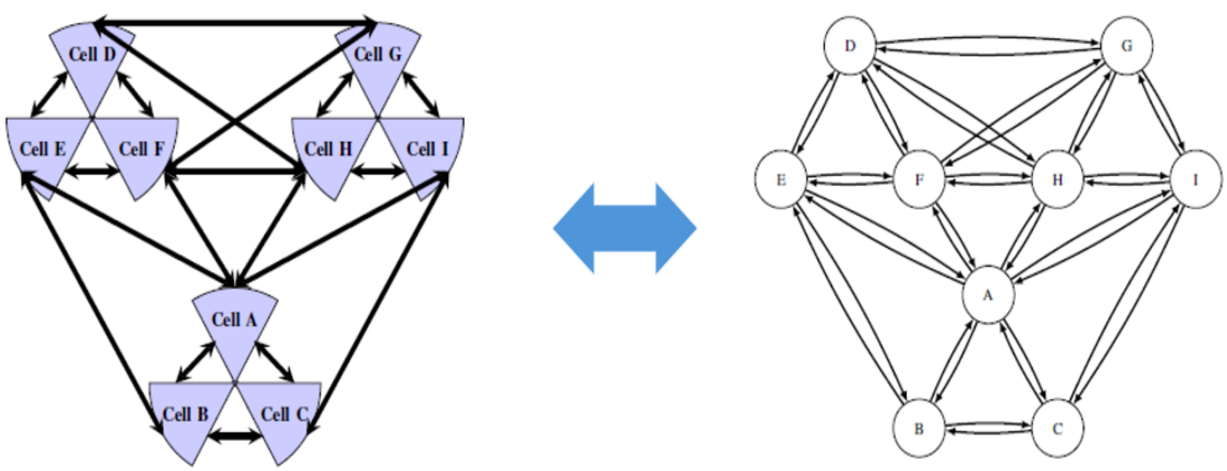

Figure 1 — Radio access networks exhibit influence-based interaction structures (left) that are structurally equivalent to interaction graphs used in neural systems (right). This correspondence motivates treating the radio network itself as a learnable system.

2.2 UE-MEDIATED OBSERVABILITY AND THE BIPARTITE NEURAL STRUCTURE

The radio access network does not observe interaction directly. Observable evidence of network behavior is mediated through user equipment. Each UE samples the local radio environment and reports measurements reflecting the superposition of signals from surrounding cells. These observations are localized, partial, and time-varying, but collectively provide complete coverage of network behavior.

The RA-NN treats the user equipment as first-class entities within the system. Each UE defines a local interaction structure over the cells it observes, capturing which cells influence its radio conditions and to what degree. This induces a bipartite structure with UEs and cells as distinct node types and UE–cell edges representing observed interaction relationships.

Users are conditionally independent given the cells they observe; coupling arises implicitly through shared cell influence and the shared radio environment. This structure preserves physical meaning while remaining scalable and forms the basis for learning network-level interaction.

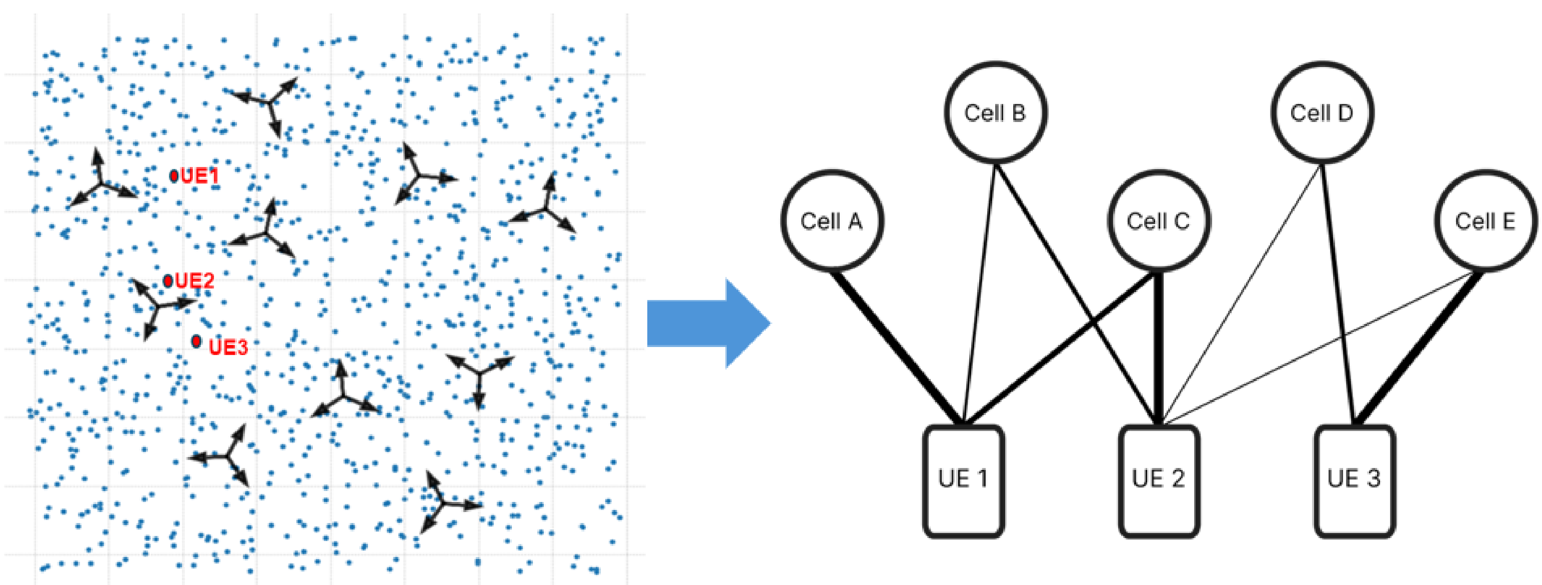

Figure 2 — Cells and user equipment form a bipartite interaction graph in which UE–cell edges represent observed radio interaction relationships derived from UE measurements. Edge weights encode the strength and context of each cell's influence on UE performance.

2.3 STRUCTURAL ALIGNMENT BETWEEN RADIO NETWORKS AND NEURAL SYSTEMS

Radio access networks and neural systems share a common structural requirement: system-level behavior emerges from the interaction of many entities whose influence on one another is neither uniform nor static. Outcomes are determined by how influence propagates through the network rather than by isolated decisions.

Neural systems make this structure explicit through learnable representations of influence. In the RA-NN, the physical overlap of radio waves is represented as intrinsic, learnable system state. This is not a metaphorical alignment; it is a direct expression of how radio networks already operate implicitly. These representations persist across control functions and learning processes, providing a stable foundation for coordinated AI-native behavior.

2.4 COVERAGE AS THE FUNDAMENTAL LEARNED SUBSTRATE

This physical cell coverage is the fundamental learned substrate underlying all radio access network behavior within the RA-NN. It is not treated as a static footprint, predefined service area, or externally computed metric. Instead, it is derived directly from user observations of the radio environment and maintained as intrinsic system state.

All radio access network behavior ultimately acts through this substrate. Service quality depends on the health of individual radio links. Mobility decisions arise when link conditions degrade and alternative links must be established. Load balancing requires physically changing radio links to cells capable of sustaining acceptable coverage. Interference management and capacity shaping operate through the same learned representation. By treating coverage as the substrate, the RA-NN unifies functions that are traditionally addressed through separate task-specific abstractions.

3. LEARNING, STRUCTURE, AND COORDINATION IN THE RA-NN

The coverage substrate in the RA-NN is induced from learned representations of individual UE–cell interactions. User equipment measurement reports provide the source of information for resolving this interaction space. Each UE contributes a localized estimate of effective link condition with respect to surrounding cells. These estimates are used to learn the radio dynamics through attention-centric modeling that assigns different embeddings to the subset of cells whose signals materially influence the UE's observed radio conditions.

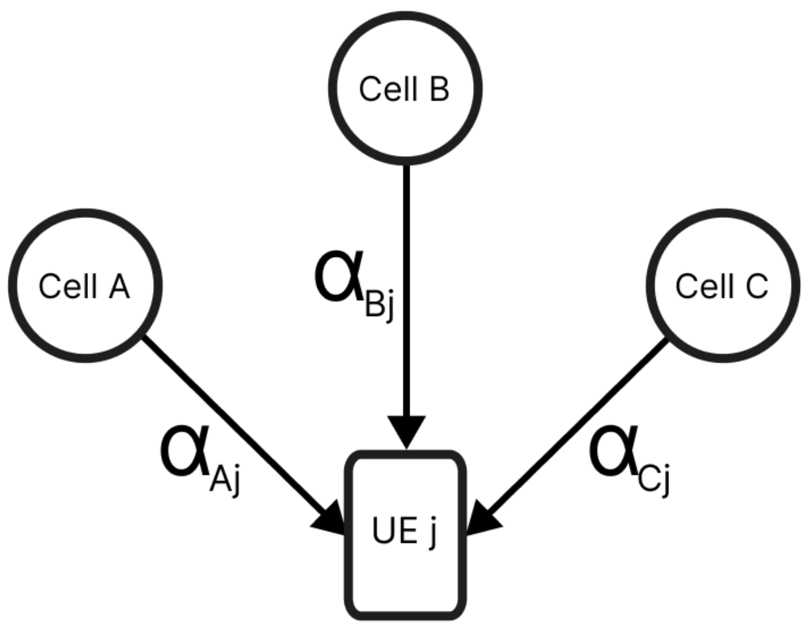

Figure 3 — Each UE assigns learned attention weights (α) to a sparse subset of surrounding cells based on observed radio conditions. These per-UE interaction weights constitute the atomic inputs to learned coverage and will happen simultaneously for all UEs in the observed network.

These per-UE links reflect the local combined effects of propagation loss, geometry, antenna gain, interference, and configuration under current conditions, without requiring explicit modeling of these components.

3.1 LEARNING THE GLOBAL INTERACTION STRUCTURE

While per-UE link state provides the atomic view of network behavior, network-level reasoning requires operating over populations of users at the cell level. Across the user population, overlapping UE observations constrain which cell–cell relationships are physically meaningful.

The RA-NN does not assume predefined topology, configured neighbor lists, or externally supplied adjacency information. No cell–cell relationships are assumed prior to observation. Instead, the global inter-cell structure is learned as part of the system's internal representation from aggregated UE-mediated observations.

This learned structure is maintained as a sparse, attention-based inter-cell graph that persists as part of the operational system state. Unlike attention used during offline model training, these attention relationships evolve continuously as the radio network is optimized and conditions change. Cells persist as stable physical entities, while inter-cell links are dynamically learned and updated.

This learned attention structure is analogous to an AI-native form of Automatic Neighbor Relations (ANR), a system that is present in legacy RAN deployments to maintain current neighbor lists to facilitate mobility.

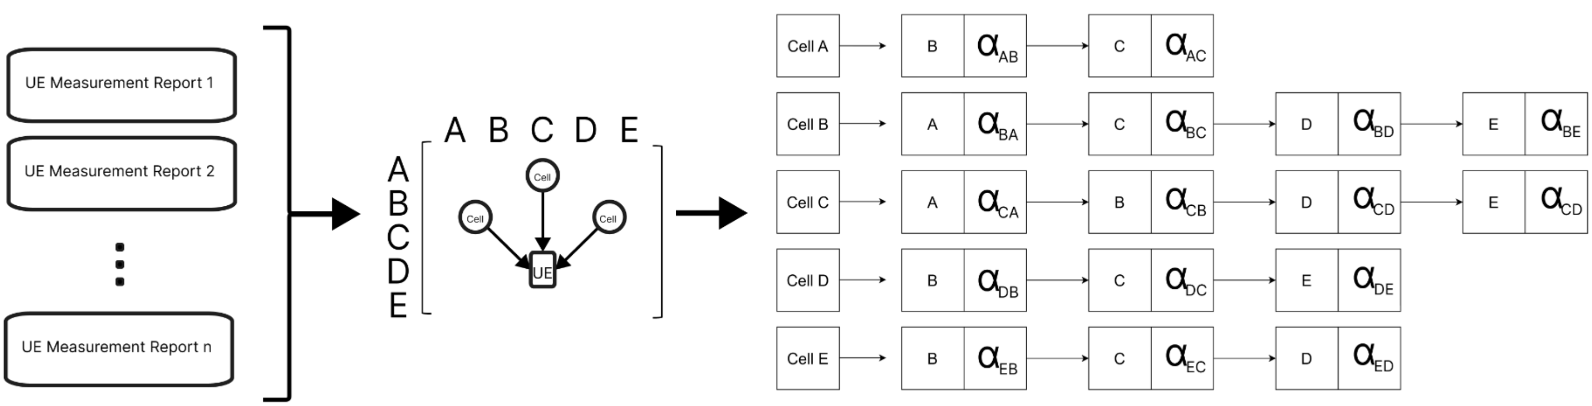

Figure 4 — Inter-cell influence is inferred from UE observations and materialized as a sparse attention-based adjacency retained as system state. The underlying interaction space is conceptually unconstrained, allowing any cell to potentially influence any other prior to learning.

By reducing the effective interaction space from a conceptual N^2 hypothesis to a learned sparse graph, the RA-NN enables scalable reasoning while remaining faithful to physical radio behavior. All coverage learning and downstream control operate within the constraints imposed by this structure.

3.2 USER EXPERIENCE AS A DISTRIBUTION OVER PER-UE LINK STATES

The RA-NN represents user experience as a distribution over learned per-UE link states rather than as cell-level counters or aggregated KPIs. Each UE remains an explicit entity within the system, contributing its own learned link state derived from observed radio conditions.

With the sparse inter-cell graph structure established, the RA-NN can reason coherently over populations of users at the network level. These per-UE states form a distribution that captures spatial variation, tail behavior, and localized degradation that would otherwise be obscured and unactionable at scale using aggregate metrics alone.

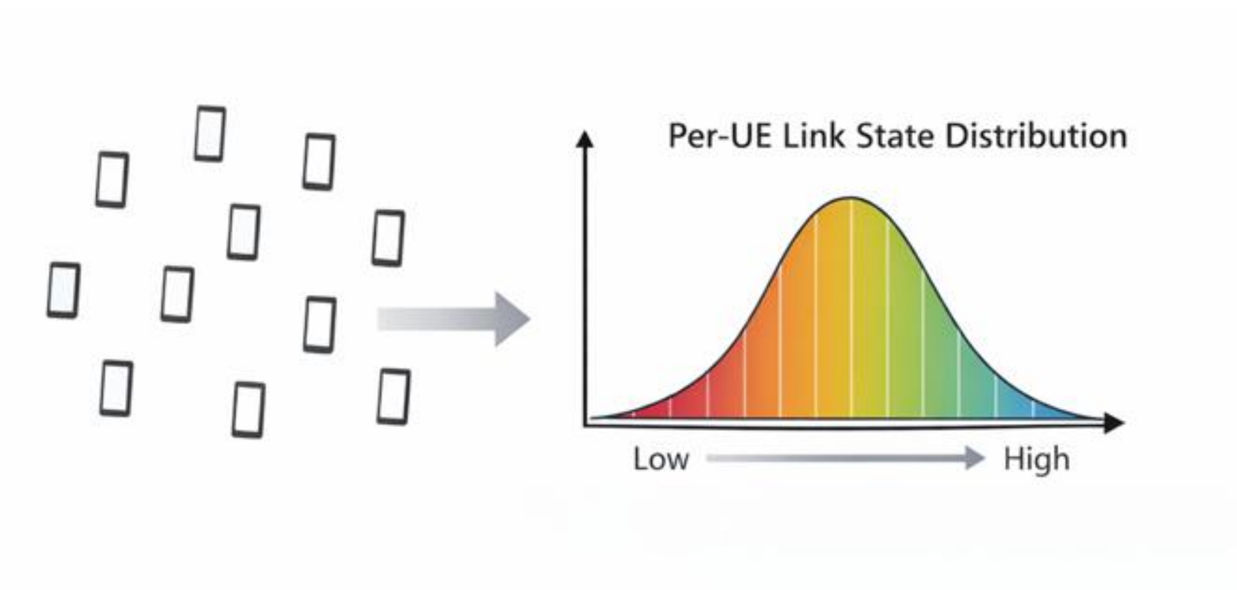

Figure 5 — User experience is represented in the RA-NN as a distribution over per-UE link states. Individual UEs contribute independent link state samples, while network-level reasoning operates over the resulting population distribution.

3.3 DOWNSTREAM COORDINATION AND MULTI-CELL CONTROL

The learned coverage substrate established within the RA-NN provides a unified basis for control across the radio access network. Coverage is represented consistently at multiple resolutions—as per-UE link state, as sparse inter-cell interaction, and as a distribution over user experience—allowing control decisions to be evaluated in terms of their impact on the underlying physical coverage of the network.

Multi cell-coordination is derived from decision-making over shared learned state, against which the cross-cell effects of actions can be assessed directly.

Control decisions may be expressed at multiple levels of granularity:

- Over subsets of the cell graph, coordinating actions across related cells to modify coverage and interaction patterns for the users they jointly influence
- Over subsets of the user population, shaping the distribution of experience for selected groups of users in different network regions
- At the level of individual users, conditioned on each UE's learned link state and local context

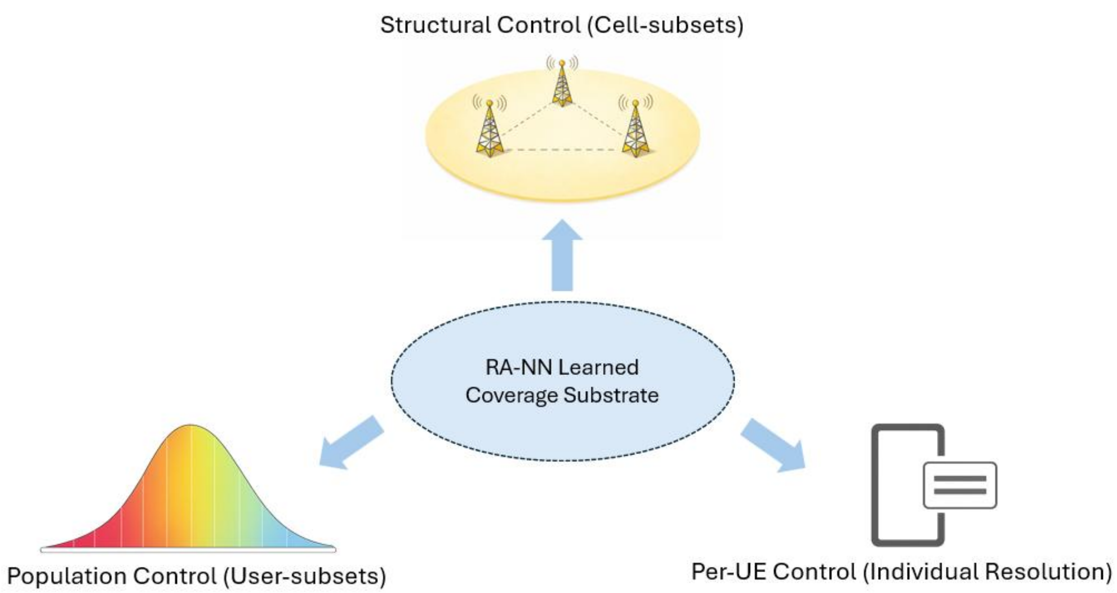

Figure 6 — The shared coverage substrate is applied across multiply granularities to provide a unified basis of coordination and control for the RA-NN

These are not separate control systems. They are different resolutions of control applied using the same learned substrate. Actions taken at one level propagate through the learned sparse interaction structure and alter conditions experienced elsewhere in the network.

3.4 LEARNING CONTROL POLICIES OVER THE COVERAGE DISTRIBUTION

The RA-NN provides a direct interface for learning how network configuration choices affect radio behavior. Actions are evaluated by their observed impact on the user experience distribution, allowing the system to associate configuration changes with their effects under real operating conditions.

The RA-NN employs reinforcement learning mechanisms to learn configuration policies. The objective is to learn a policy that maps observed network state to configuration choices based on their expected effect on coverage and user experience. Configuration choices that move the user experience distribution in desired directions are reinforced.

Learning signals are obtained directly from the radio environment. Unlike supervised learning, where gradients are derived from labeled targets, the RA-NN derives policy gradients from observed changes in coverage, interference, and user experience following configuration updates. In this formulation, the physical network provides the supervision required to learn the effects of configuration changes.

Reward functions are defined in domain-specific terms to encode what constitutes improvement in a radio network, ensuring that learning aligns with physically meaningful and operationally valid objectives.

Because learning operates on the shared coverage substrate, policy learning is defined uniformly across control resolutions. The RA-NN learns how configuration actions modify coverage and interaction structure across cells and users as an intrinsic property of the system, rather than as a collection of task-specific optimization loops.

APPENDIX A: RA-NN FORMULATION DIAGRAM

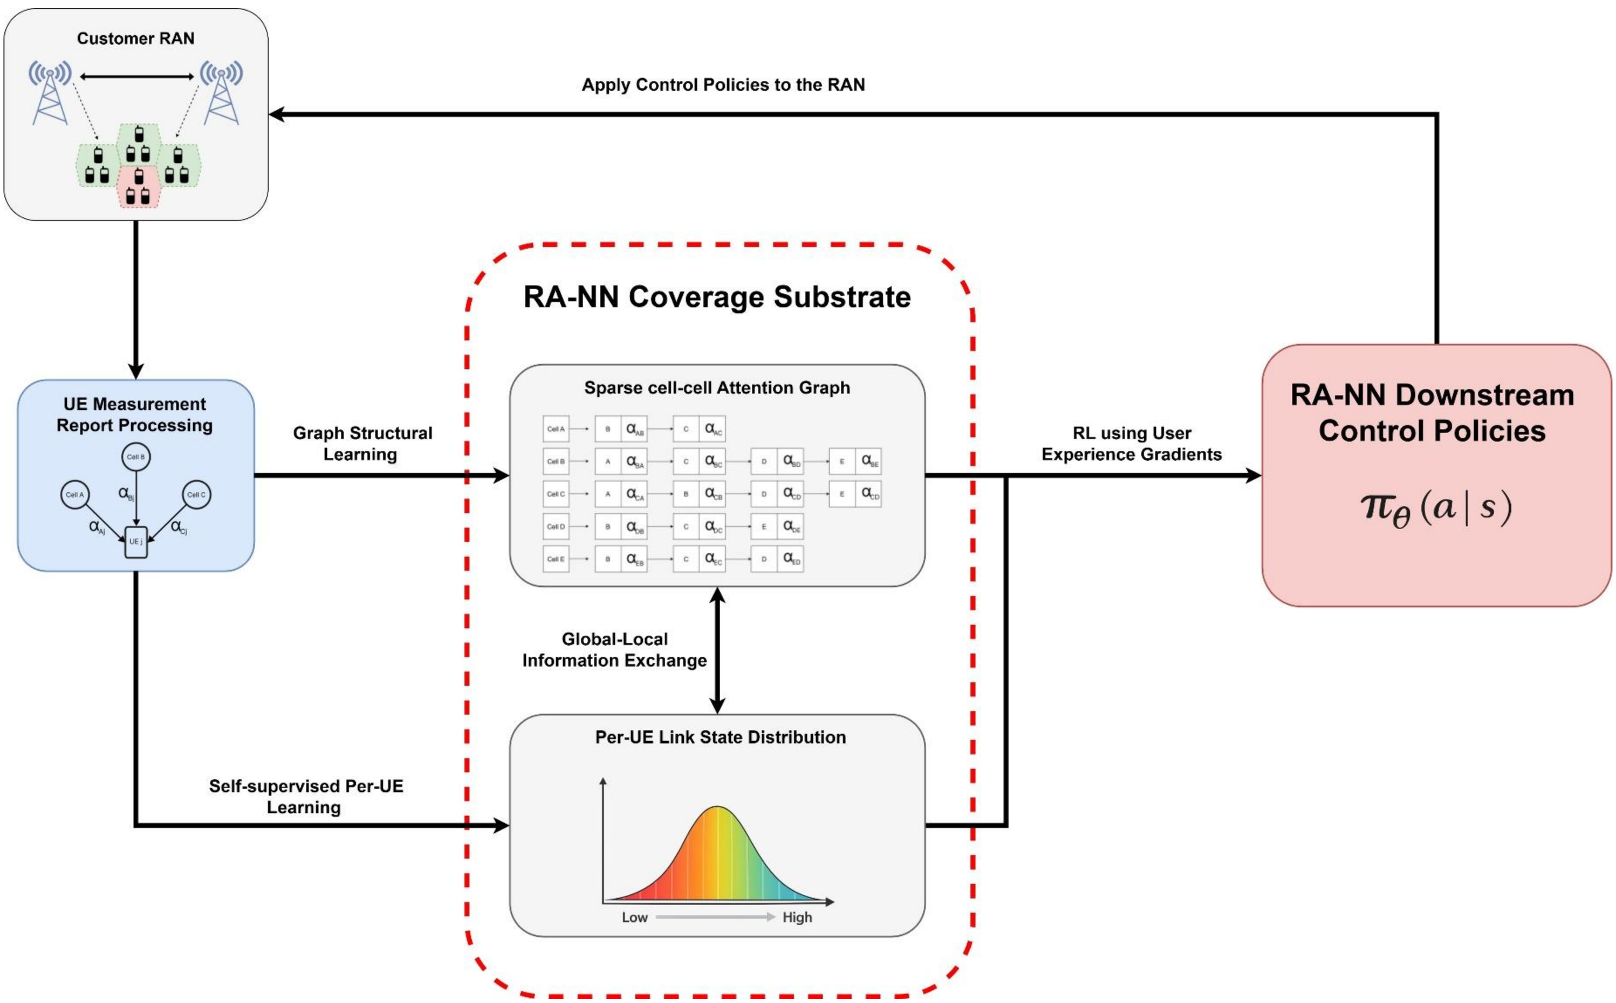

Figure 7 — Full RA-NN system diagram showing the AI-native coverage substrate that learns a probabilistic cell–cell topology and per-UE link state distributions from UE measurements. These learned representations provide the foundation for downstream control policies $\pi\theta(a|s)$, optimized via user experience feedback and applied directly to the RAN.