Advancing Spectrum Science: Rely on Domain Knowledge or Handoff to Machine Learning Models?

Fatemeh Afghah
Associate Professor, Department of Electrical and Computer Engineering
Director, Intelligent Systems and Wireless Networking (IS-WiN) Lab.
Clemson University

CLEMSON UNIVERSITY

Sep. 2025

INTELLIGENT SYSTEMS AND WIRELESS NETWORKING LAB
CLEMSON UNIVERSITY

Email: fafghah@clemson.edu, IS-WiN Lab: lhttps://iswin-lab-clemson.github.io/Home.html

AI vs Domain Knowledge in Communication systems

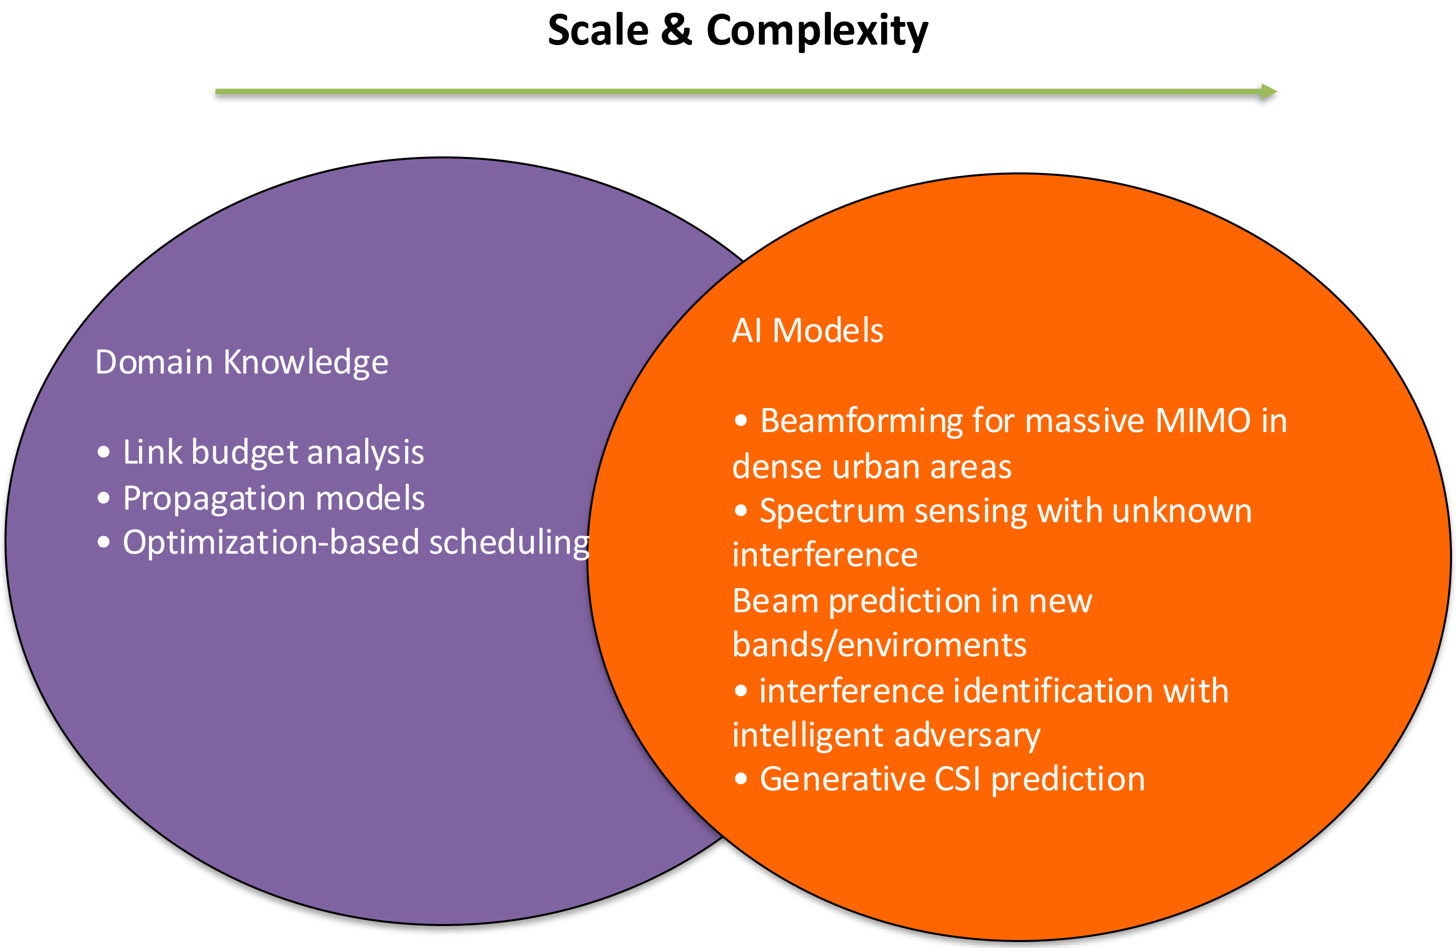

Hybrid (Gray Zone)

- Physics-Informed Neural Networks
- AI-augmented channel estimation
- Human-in-the-Loop Reinforcement Learning

From Images and Robots to Wireless: What Tasks Shaped Foundation Models?

NLP	CV	Robotics	Communications
• General-purpose reasoning • Summarization, QA, dialog • Few-shot and zero-shot generalization • Agent frameworks	• Joint embedding with text (CLIP) • Image retrieval & recognition • Visual Question Answering (VQA) • Segmentation (object boundaries, SAM)	• Object detection & pose estimation • Waypoint planning in cluttered environments • Skill learning & transfer across tasks • Embodied VLMs (language → action)	• Spectrum sensing & interference ID (below-noise detection, real-time classification) • Cross-layer multi-objective optimization (RL) • Network digital twins • Beam management in massive MIMO
Led to GPT, Gemini, LLama	Led to ViT, CLIP, BLIP, SAM	GENESIS,Gato, Dearmer, Muze	**???**

CV/NLP/Robotics → **labels and tasks** drove foundation models.

In Communications → **the physics itself is the task**. Needing models that respect **causality, propagation laws, and real-time constraints**.

Pretraining Foundation Models for Wireless

- **Physics is the data generator.** We don't have semantic labels for "channel images," but we *do* have **exact constraints** (superposition, reciprocity, passivity, sparsity in angle–delay, WSSUS, path-loss scaling, Doppler structure).
- **Hard real-time and control coupling.** Inference must feed decisions (beam, scheduling, handover, RIS phase) at sub-ms scales.
- **Multi-view measurements, not captions.** We often have *paired but heterogeneous* readings: UL/DL reciprocity, adjacent subcarriers/times, neighboring locations, BS/UE antenna domains, ray-tracer vs. measured traces(Time–frequency–space masking , Neighbor contrastive, Sim-to-Real alignment)
- **Cross-modal grounding (RF × Vision/LiDAR maps).** align channel embeddings with **scene geometry** (satellite/StreetView/LiDAR) via contrastive loss

Challenges toward Wireless LM

- **Scale vs fidelity**: LLMs are huge and trained on text; embedding physical laws explicitly can be hard and computationally demanding.
- **Simulators in the loop**: Using external simulators helps, but can limit generality and speed; approximations may degrade with more complex systems.
- **Correctness / generalization**: Even when models are physics-informed, their physical reasoning can break in edge cases or where the physics is subtle or has many interacting components.
- **Training data**: Many physics laws are not well represented in text or even multimodal datasets; obtaining appropriately labeled or simulated data is expensive.

Beyond Communication: Joint Planning & Sensing for Autonomous Swarms in Dynamic Environments

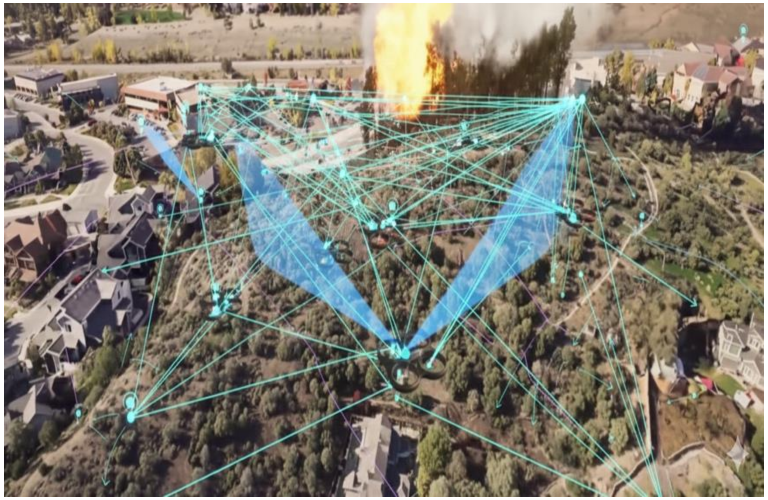

Lack of Prior Information: lack a priori knowledge of new signal types (e.g., new threats)

Lack of Stationarity: complex non-stationary environments

Coupled Problems: Communication is no longer a standalone task. It is intrinsically linked with sensing (e.g., locating a survivor) and planning (e.g., safely routing a UAV swarm). These interconnected problems explode in complexity.

ORAN

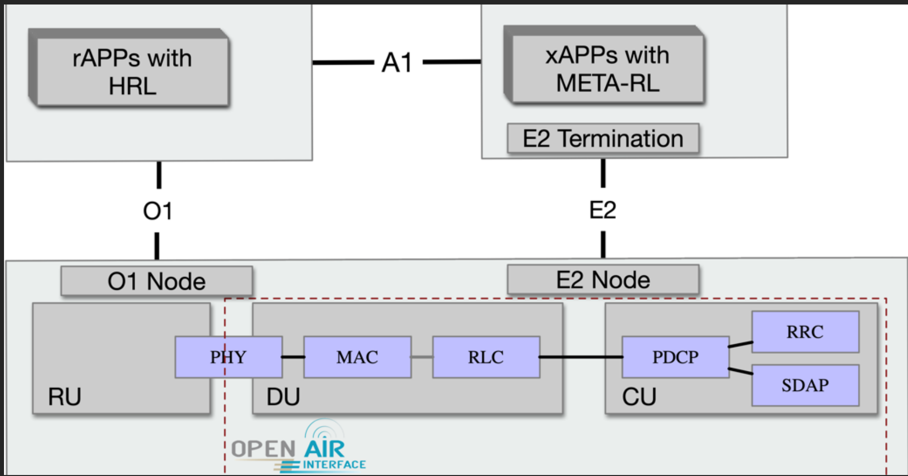

- **Domain prompts (expert context):** ORANSight reads live RAN stats and emits short text prompts that *explain* the situation.
- **Learnable prompts (tuning the LLM):** A small set of trainable “soft tokens” is prepended to the domain prompt. These tokens are optimized during training so the frozen GPT turns the mixed input (raw state + domain text) into embeddings that highlight what matters *for control*.
- **Knowledge distillation (lightweight deployment):** During training, the student (GPT+learnable prompts) is forced to **mimic** the teacher (ORANSight) via a distillation loss.

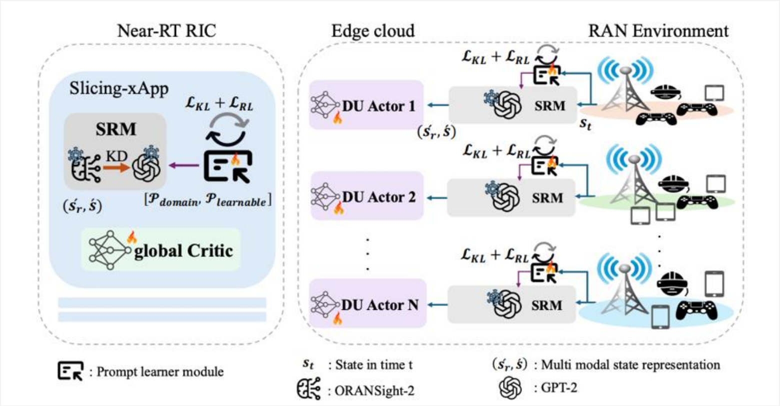

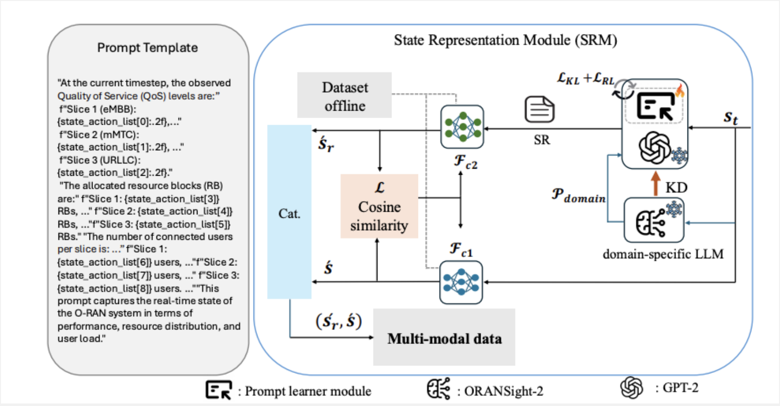

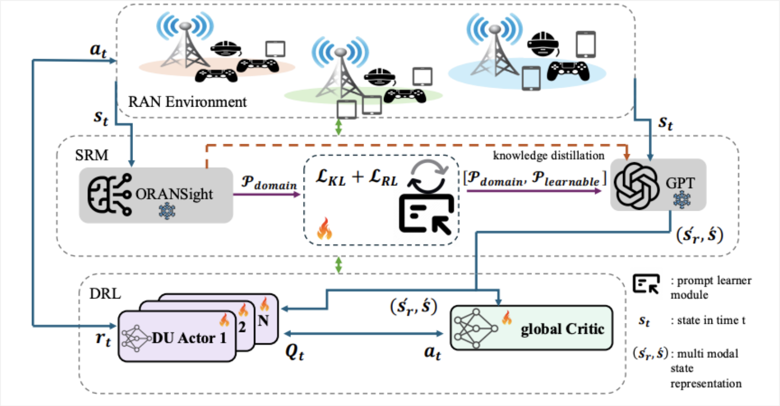

ORAN-GUIDE: RAG-Driven Prompt Learning for LLM-Augmented Reinforcement Learning in O-RAN Network Slicing, 2025.
LLM-Augmented Deep Reinforcement Learning for Dynamic O-RAN Network Slicing, ICC 2025.