From AI in the RAN to AI for the RAN: Making innovation safe

Striking a balance between benefiting from AI and protecting networks, users and society from its excesses. A conversation with Nirlay Kundu, Head of Technology Standards at IMDEA Networks, and Monica Paolini, Principal at Senza Fili

SENZA FILI

From AI in the RAN to AI for the RAN:

Making innovation safe

Striking a balance between benefiting from AI and protecting networks, users, and society from its excesses. A conversation with Nirlay Kundu, Head of Technology Standards at IMDEA Networks, and Monica Paolini, Principal at Senza Fili

Monica: Today, we are going to talk about how AI is changing the RAN with Nirlay Kundu at IMDEA Networks in Spain. Nirlay has been working on the evolution of the RAN for a long time and, before joining IMDEA Networks, he was part of Verizon's Open RAN standardization efforts and held a leadership role in the O-RAN ALLIANCE.

Nirlay, how has your involvement in the RAN transformation changed over the years?

Nirlay: As I moved from Verizon to IMDEA Networks, I continued my standards work, and I represent IMDEA Networks in the O-RAN ALLIANCE. I also help research students at the institute in a technology-agnostic role, working on networking interfaces such as 6G or 5G, as well as on technologies such as AI and blockchain.

Can you tell us more about IMDEA Networks?

Nirlay: IMDEA Networks Institute is a research organization on data networks whose multinational team is engaged in cutting-edge fundamental science and technology. As an English-speaking institute in Madrid, Spain, IMDEA Networks offers a unique opportunity for pioneering scientists to develop their research vision. IMDEA Networks has established itself internationally at the forefront in the development of future network principles and technologies. Our team of highly reputed researchers is designing and creating today the networks of tomorrow.

Monica: We all agree that AI is going to have a major impact on the RAN. But will it abruptly redefine the RAN, or will it foster a gradual transformation?

Nirlay: Today, we are deploying AI in the RAN. Most AI deployments in telecom are geared toward optimizing existing architecture. Think of AI as a technician that tunes parameters, detects anomalies, schedules resources, and predicts traffic. These are often narrowly scoped and data-driven tasks, bound by existing protocols and infrastructure. This is smart automation, not strategy transformation.

What is emerging now is AI for the RAN. This is where the conversation becomes more transformative. We are starting to ask what a

The role of AI in the RAN is evolving from AI in the RAN, to incrementally optimizing RAN operations and increasing cost efficiency, to AI for the RAN, to reimagine the RAN architecture, ensuring accountability and resiliency and protecting networks, users, and society from biases and unintended outcomes.

Nirlay Kundu, Head of Technology Standards at IMDEA Networks, guides us through a strategic approach to benefiting from and safeguarding against AI, drawing on his research and standardization experience, as well as his previous work at Verizon.

We talked about:

- **Accountability and governance with AI**
- **Predictive, adaptive, autonomous and explainable intelligence**
- **The essential role of humans in the loop**
- **Minimizing bias, maximizing learning**
- **AI at the edge and in the data center**
- **The latest innovations in agentic AI**

Watch the video of the conversation

RAN should look like if AI is part of the design process from day one. This includes AI-informed spectrum policy and spectrum-sharing models, AI-guided slicing based on societal, economic, and mission-critical applications, and AI-driven architectural choices, such as when to deploy microcells, when to use satellite, and when to use edge compute. AI can also help rethink coverage, investment decisions, and even fairness using intelligent simulations and reward models.

AI for the RAN is not just about efficiency. It is about changing the goals, not just achieving them faster. Today's AI-in-the-RAN discussion is mostly about vendor capabilities, while the emerging AI-for-the-RAN discussion focuses on high-level regulatory, boardroom, policy, and strategy considerations. If we do not ask AI to help us redefine what we want the RAN to do, including societal goals such as digital inclusion, energy justice, and public service access, we risk building faster networks that replicate past inequalities. That is why it is important to discuss AI for the RAN now.

Monica: We often think about using AI to enable new applications and services, especially for the enterprise. With AI for the RAN, the network can fundamentally transform itself and become more closely tied to physical, social, political, and economic environments, increasing its usefulness and fairness.

But what kind of intelligence are we looking at for the RAN? Is artificial general intelligence (AGI) going to be relevant for the RAN if and when it becomes available?

Nirlay: Your question gets to the core of what we mean by AI for the RAN. The term "intelligence" is often used loosely in this context. For clarity and strategy planning, it is essential to understand what kind of intelligence we are dealing with. We should think of RAN intelligence not as a single capability, but as a spectrum of functions, each with different implications for design, control, accountability, and risk.

AI gives us predictive intelligence. It anticipates future states using time-series, historical, or real-time data. Example use cases include predicting traffic surges, anticipating handovers, and forecasting peak demand. Instead of reacting to events after they happen, operators can adopt a proactive approach. We still need human intervention and validation because outcomes depend on data quality and generalization error.

In the next phase, autonomous intelligence allows AI to make and execute decisions without human intervention. Use cases include self-tuning of antenna parameters, automated load distribution, and closed-loop network optimization. AI models can act faster than human operators, scale more easily, and reduce operating expenses by removing manual intervention.

The challenge is that AI can be difficult to govern. Because decision-making is often opaque, there is a risk of cascading failures. We still need to determine who is responsible if autonomy causes a service disruption or complication.

Another form of intelligence is adaptive intelligence, in which behavior evolves in response to changing conditions through a feedback loop. It enables networks to adapt to environment-specific usage shifts, new devices, and new mobility patterns, reducing model retraining costs and improving robustness over time. However, adaptive intelligence can drift in unintended ways, requiring robust monitoring and rollback mechanisms.

We also have explainable intelligence, which provides a traceable and understandable rationale for its actions and outputs. How do you decide how to allocate spectrum? When do you generate a dropped-connection alert? These capabilities are critical for compliance, debugging, and trust-building, especially with regulators and customers. The trade-off is between model complexity and performance.

Most AI deployed today is predictive and semi-autonomous. We are rapidly moving toward systems that are adaptive and increasingly opaque, which makes explainability and accountability non-negotiable.

Monica: AI is not just about crunching more data; it can also identify patterns that humans miss because of their experience bias. If something works today, I am going to do the same thing tomorrow. AI continues to explore alternative approaches across domains and use cases. How should we manage this AI-driven continuous exploration?

Nirlay: There is a learning curve. You need to manage interactions between autonomous systems and humans. A Responsible, Accountable, Consulted, Informed (RACI) matrix is required to define who is responsible for what. This is an extension of software entities that learn autonomously. There should be a feedback loop with continuous evaluation, in which humans assess correctness and intervene to adjust and improve algorithms when outcomes are misaligned with objectives.

Monica: We need to recognize that AI can bring benefits but also cause harm. How do we

© 2026 Senza Fili

manage that, and who takes responsibility? For example, autonomous driving can cause accidents, but the accident rate may be lower than with human driving. Still, people tend to be more critical of mistakes in AI-driven decisions or services, and we need to develop a more deeply grounded approach.

Nirlay: This is a critical issue. While AI promises greater efficiency, it introduces complexity, accountability challenges, operational risk, and legal issues. Ultimately, consumers use operator services, and operators own the customer experience, network performance, and quality of experience. Even if operators outsource decision-making to AI systems or third-party vendors, they are ultimately legally and contractually responsible for Service Level Agreement (SLA) compliance, Quality of Service (QoS) and Quality of Experience (QoE) guarantees, and regulatory compliance for things such as lawful intercept, net neutrality, and spectrum usage.

Right now, if AI changes something in the network, the operator is accountable. For example, if an operator reduces transmission power to save energy, it is responsible for any resulting service degradation, even if the decision was made by a vendor's AI. Operators must decide whether to deploy AI recommendations automatically or with a human in the loop, ensuring AI outputs are interpretable and traceable within the network operations center. Operators need a monitoring model to detect drift or abnormal behavior after deployment. Even when vendors supply AI, operators decide how and where they apply it and are liable for the outcome.

Governance, regulatory compliance, and risk management are also critical. Operators must maintain AI risk logs, implement fallback systems, and retain the ability to revert to legacy or rule-based optimization when things go wrong. We need to establish compliance mapping between AI decisions and telco regulations, and to be able to audit AI actions for forensic purposes. For example, we need to find out why a service outage or a spectrum violation happened: Was it AI or was the human in the loop? Regardless of the systems used, operators remain responsible for customer harm, such as unfair throttling, uneven access, or dropped calls.

Vendors, by contrast, provide the AI software and hardware and develop the algorithms, models, and platforms that inform and automate RAN control decisions. Unless specified in contracts, vendors are not legally liable for operational outcomes. They are responsible for the model architecture, the model training pipeline, validation in the test environment, such as simulation labs, and synthetic data. They track performance metrics, such as accuracy, convergence speed, and impacts on KPIs, and must satisfy SLAs. But this is where their responsibility stops. Models are trained on general datasets. Vendors do not have, or have only partial, real-world datasets from the operator, and cannot conduct full-fledged testing because operators will not share much of their data. This can lead to a generalized approach from vendors.

When vendors cause an AI-driven failure, they should document the failure modes and limits of model performance. They should clarify what data the AI model expects and how to avoid garbage-in-garbage-out outcomes. They should provide some error margins at regular intervals, specifically for high-stake decisions, for example, for spectrum assignments or handovers. Vendors are responsible for the overall security and robustness of the model and for ensuring that it is not vulnerable to manipulation, intentional or unintentional. They need to protect the model's intellectual property and data pipelines against attacks, certify compliance with standards such as those from the International Organization for Standardization (ISO) and the O-RAN ALLIANCE, and follow security best practices. That is the vendor's responsibility.

Monica: You raise crucial issues on how using AI creates the need for a new approach to assigning responsibility when things go wrong. But can we use AI to check AI? And maybe we will have to do that because the models are too complex for humans to analyze.

Nirlay: We are just at the beginning. A slight distortion of the data yields synthetic data for model training, initially resulting in a minor aberration. But when you run it in a loop 100 times or 1000 times, it just goes wild. That is what I am afraid of: how much real data we have and how much is the delta per loop. It ends up with a nasty prediction. This makes me worry about the capacity of AI monitoring AI.

Monica: This is similar to what happens with GenAI in the wider domain, beyond telecoms. LLMs are regurgitating what is already out there. If what is out there is false or inaccurate, AI amplifies the distortion. If you feed an AI model to optimize the network with inaccurate data, it is likely to go down the wrong tangent and may or may not be able to self-correct. To be fair, this happens to humans, too, but the risk is on a different scale with AI.

This raises the question of how we should invest in and adopt AI. Not only do operators need to ensure everything works as planned, but they

also need to decide where and when to integrate AI into the network. Where should operators deploy AI, at the edge or in the cloud? Are they overinvesting or being too slow? How much computational power do they need at the edge? Where do they need GPUs?

Nirlay: My advice is to start slow and small, and see whether AI produces results, whether deployed at the edge or in a data center. Recent developments show we can achieve high-performance results without high-power GPUs through algorithm optimization. Software performance is going to improve quickly right now because everyone is thinking about optimizing their software and using less GPU power, compute, and energy.

At the end, you need to figure out how much you are saving. Are you saving on opex? Are you saving on capex? Mostly, it will be opex savings. If the savings are on the bottom line, or if the numbers are bigger on the top line, that is where we should focus. Where is the biggest bang for the buck? We are still learning in this game. Focus on the areas with the greatest return on investment, and avoid concentrating all efforts in a single area. Experiment across different parts of the network to identify where you create the most value.

Monica: Not all operators will reach the same conclusion. A common concern is that AI requires so much power that the higher energy costs offset any opex savings from increased network efficiency. I would argue that as models improve, training and inference will become more energy efficient.

Nirlay: Yes, this is a long-term effort. Another way to address this is to use SLMs, which are less energy hungry and can address some privacy issues, as the data can be confined to within the company premises.

Monica: This creates challenges because in telecoms, as well as in other sectors, we look at the short-term financial impact. We look at quarterly results, and AI is not a quarterly kind of thing.

Another potential benefit of AI is its ability to discover things we did not know. But because we did not know them, it is difficult to assess their value. What is the best way to transfer our expertise, accumulated over a long period across multiple domains, to an AI model?

Nirlay: This is a crucial question, especially in the context of AI for the RAN, where decisions can affect millions of users, critical infrastructure, and regulatory compliance. AI lacks visibility into many physical realities of the RAN, i.e., what AI can't see and what sensors don't capture, such as local topography shifts, new buildings that block signals, microclimate or seasonal changes, and seasonal interference. Let's say a field technician's actions are not logged in, or that hardware faults or fiber cuts are not reported. AI doesn't know exactly what is going on. The AI may optimize based on an incomplete or outdated view of reality. That is an urgent issue right now.

AI does not understand human context and intent. AI doesn't understand why we set the policies or thresholds we do. For example, do AI models understand why emergency coverage requirements for hospitals are as they are now? Do they understand the legal and political implications of prioritizing one region over another? AI still doesn't understand the context well. Socioeconomic factors are also a big challenge because they are subjective, for example, in spectrum allocation. AI may optimize technically, but it may do so unethically or illegally because of the lack of understanding of the wider socio-economic human context.

Another concern is that AI is trained on time-series past data and may fail in new scenarios. Suddenly, there are mass gatherings, protests, or festivals. Out of the blue, something is happening. Let's say there are cyberattacks that cause interference. Recently, in Spain, there was a power outage. If AI had acted on that, it would have been a failure. Humans intervened to take the situation under control. AI is not intelligent enough today to respond to malicious interference and cyberattacks. With drones and autonomous cars, we may suddenly see congestion that AI cannot predict. AI would not know why new mobility patterns, generated by devices it does not know about, happen. The AI may not know how to handle these events and may make the wrong decisions. Humans are still required to intervene in such cases.

The training data fed to AI models is biased. Under-representation of certain communities, for example, low-income or rural communities, in the training data may be unintentional but have undesired results. Vendor data silos across the RAN and the core, clerical mistakes, and manual errors may amplify AI biases, leading to uneven performance or unintended outcomes.

To address these limitations, a human-in-the-loop approach remains essential. AI systems must allow expert override in high-impact situations. We need a backup plan for human intervention to enforce quality of service (QoS). We need to develop confidence in the model to resolve this uncertainty.

AI should express its confidence. Maybe there

could be a flag for out-of-distribution scenarios and support rollback mechanisms. Every autonomous loop should include clear, well-documented, safe-mode mechanisms that revert to conservative, well-tested methodologies.

AI in the RAN is powerful, but not ubiquitous. It still has its blind spots, and we cannot ignore them. They must be anticipated, monitored, and governed. Good AI does not assume that you don't know everything. The model should not be the be-all and end-all. Ideally, AI models should recognize their limits and request human intervention when they encounter something they do not know.

Monica: AI models need guidance, just as kids need education. Just as you do with kids, you need to push models in the right direction. Humans must provide strategy and direction to prevent something bad from happening. For instance, you may want to guide AI models to recognize and address bias in the data, for instance, to ensure you do not worsen the digital divide as we discussed earlier.

Do you know when you are not overcompensating for the bias in the data? Are we substituting the initial bias with another bias? How is this different from what we are already doing to correct for bias in our assumptions and data? AI amplifies bias, and we need transparency to be aware of and minimize its impact. It is a challenge to devise where to draw the line because there are security and ethical tradeoffs. For instance, do you balance security with a potential reduction in the quality of experience?

How does agentic AI fit into all this?

Nirlay: Agentic AI is a crucial area of AI development today because it enhances the network's autonomy. It operates independently, deciding and acting on its own without constant human guidance. It is able to learn and adapt to new information in collaboration with humans. It can provide assistance and insight into complex situations. By relying on pattern recognition and probabilistic reasoning, it provides flexibility and adaptability in dynamic conditions.

About IMDEA Networks

IMDEA Networks Institute is a research organization on data networks whose multinational team is engaged in cutting-edge fundamental science and technology. As an English-speaking institute in Madrid, Spain, IMDEA Networks offers a unique opportunity for pioneering scientists to develop their research vision. IMDEA Networks has established itself internationally at the forefront in the development of future network principles and technologies. Our team of highly reputed researchers is designing and creating today the networks of tomorrow.

Some keywords that define us: 5G, 6G, Big Data, blockchains and distributed ledgers, cloud computing, content-delivery networks, data analytics, energy-efficient networks, fog and edge computing, indoor positioning, Internet of Things (IoT), machine learning, millimeter-wave communication, mobile computing, network economics, network measurements, network security, networked systems, network protocols and algorithms, network virtualization, software defined network (SDN), network function virtualization (NFV), privacy, quantum communication, social networks, vehicular networks, wireless networks and more.

About Nirlay Kundu

Nirlay Kundu is a telecom thought leader with 25+ years of experience spanning global operators, vendors, startups, and academia. He has led strategic partnerships, standards initiatives, and innovation programs across 5G, Open RAN, cloud, and emerging 6G technologies. A former Distinguished Engineer at Verizon, he has worked extensively with Tier-1 operators and standards bodies, including the O-RAN ALLIANCE, ITU-T, and 3GPP. Nirlay currently bridges industry and academia as Head of Technology Standards at IMDEA Networks and Professor of Practice at leading universities. He holds an MBA from Babson College and an MS in Engineering from the University of Leicester.

About Senza Fili

Senza Fili provides advisory support on wireless technologies and services. At Senza Fili, we have in-depth expertise in financial modeling, market forecasts and research, strategy, business plan support, and due diligence. Our client base is international and spans the entire value chain: clients include wireline, fixed wireless, and mobile operators, enterprises and other vertical players, vendors, system integrators, investors, regulators, and industry associations. We provide a bridge between technologies and services, helping our clients assess established and emerging technologies, use these technologies to support new or existing services, and build solid, profitable business models. Independent advice, a strong quantitative orientation, and an international perspective are the hallmarks of our work. For additional information, visit www.senzafili.com.

About Monica Paolini

Monica Paolini, PhD, founded Senza Fili in 2003. She is an expert in wireless technologies. She has helped clients worldwide understand technology and customer requirements, evaluate business plan opportunities, market their services and products, and estimate the market size and revenue opportunity of new and established wireless technologies. She frequently presents at conferences and other industry meetings, and has written many reports and articles on wireless technologies and services. She has a Ph.D. in cognitive science from the University of California, San Diego (US), an MBA from the University of Oxford (UK), and a BA/MA in philosophy from the University of Bologna (Italy).

© 2026 Senza Fili. All rights reserved. The views and statements expressed in this report are those of Senza Fili, and they should not be inferred to reflect the position of the sponsors or other parties involved in the report. The document can be distributed only in its integral form and acknowledging the source. No selection of this material may be copied, photocopied, or duplicated in any form, by any means, or redistributed without express written permission from Senza Fili. While the document is based on information that we consider accurate and reliable, Senza Fili makes no warranty, express or implied, as to the accuracy of the information in this document. Senza Fili assumes no liability for any damage or loss arising from reliance on this information. Trademarks mentioned in this document are the property of their respective owners. Graphics by Senza Fili, Gemini.