Essay

Not peer-reviewed version

Automating Wireless Research and Development: What Remains Human?

Emil Björnson * , Mischa Dohler , Jakob Hoydis , Robert W. Heath Jr.

Posted Date: 16 April 2026

doi: 10.20944/preprints202604.1172.v1

Keywords: artificial intelligence; wireless communications; AI-assisted research; research automation; human–AI collaboration; agentic systems; digital twins; engineering education; technical publication

Essay

Automating Wireless Research and Development: What Remains Human?

Emil Björnson [1,*], Mischa Dohler [2], Jakob Hoydis [3] and Robert W. Heath Jr. [4]

[1] KTH Royal Institute of Technology
[2] Advanced Technology Group, Ericsson Inc., Santa Clara
[3] NVIDIA
[4] University of California, San Diego
* Correspondence: emilbjo@kth.se

Abstract

The rapid advancement of artificial intelligence (AI) is fundamentally disrupting research and engineering. While much attention is given to how AI may optimize wireless systems, this article explores a different question: how will AI impact the ecosystem and community developing future wireless technology? We trace this transformation across the entire lifecycle, from education and core research to technical publication and production-ready network deployments. As AI increasingly automates routine tasks, the primary value of the human researcher will shift from problem-solving to problem-finding, research orchestration, and oversight of trade-off management. By actively preserving spaces for deep, unplugged thinking and steering AI toward genuine discovery rather than mere recombination, we can navigate this profound shift to ultimately elevate human ingenuity and the future evolution of the researcher.

Keywords: artificial intelligence; wireless communications; AI-assisted research; research automation; human–AI collaboration; agentic systems; digital twins; engineering education; technical publication

---

Introduction

The success of distinguished researchers and engineers does not follow a single formula. As a natural consequence of finite time and capacity, individuals develop unique profiles: while some experts rely on deep knowledge in a specific niche, others excel by bridging disparate fields. Some are skilled at identifying critical problems, while others thrive at solving them. Some researchers are exceptional teachers, writers, or presenters, while others are visionaries who can inspire large communities. Today, however, with the rapid rise of AI, we witness the increasing automation of these traditionally human capabilities.

Over the last few years, we have seen frontier AI models match or surpass most human experts on a growing range of demanding benchmarks in diverse domains. Such systems have processed a vast corpus of human knowledge and demonstrate expert-level proficiency not just in a single domain, but across all fields simultaneously. They appear to solve many technical problems faster than any human and tailor explanations to diverse audiences. At the same time, the power efficiency and inference speed of AI models are expected to improve dramatically in the next few years, which will make frontier AI widely accessible at reduced cost [1].

In light of these developments, it is natural to wonder about the future role of humans in research, education, and engineering. What are the defining traits of successful researchers of the future? What should we teach, and how? Which skills will remain relevant, and which new competencies will be required? How far can we automate the research process, and how should we manage AI-generated inventions? Where will we find professional fulfillment and purpose if AI handles the heavy lifting? These are fundamental questions without clear answers. In this article, we focus on wireless communications and offer perspectives and predictions drawn from our experience in academia and

 doi:10.20944/preprints202604.1172.v1

industry. This article is not about how AI can enhance wireless systems, but rather how AI will impact the stakeholders developing future wireless technology.

To explore the above questions comprehensively, the structure of this article follows the natural progression of technological innovation—from the cultivation of talent to the deployment of real-world solutions. We begin by examining the foundational stage of education and upskilling in the AI era. Next, we discuss the core research phase, where new concepts are generated and evaluated. We will then explore how AI is reshaping technical publications and knowledge dissemination, and investigate the transition from theoretical ideas to production-ready developments. We conclude the article with our final thoughts on what lies ahead for our research and innovation community.

Future Education and Upskilling

The educational system must evolve whenever new tools for work automation become widely available. When work tasks previously performed by junior employees are automated, the next generation of graduates must be able to start working at a more advanced level than the previous generation [2]. It is a major pedagogical challenge to teach students to use advanced tools while maintaining human agency and preventing knowledge erosion [3], but we have managed to do so before. Kids learn to perform arithmetic before relying on calculators, and to write by hand before using software that corrects spelling and grammar. Similarly, a foundational understanding of advanced university topics, such as wireless communications, must precede AI automation. The advent of large-scale AI models challenges traditional teaching practices by offering personalized tutoring at a scale economically infeasible with human tutors. Over the past decade, related efforts have been explored within the blended learning framework [4], which seeks to move beyond traditional unidirectional lecturing in classrooms by shifting basic content delivery to digital platforms, freeing in-person classroom time for interactive discussions, problem solving, and laboratory activities. It is challenging to find the right setup for such practices, but the rapid maturation of AI-mediated learning tools will likely accelerate this paradigm shift for better or worse.

Students of wireless communications should not be trained to memorize theorems and formulas, but to know their existence and understand their meaning, scope of applicability, and inherent limitations. These skills are related to the three lowest cognitive levels in Bloom's classical taxonomy [5], illustrated in Figure 1. Future education must be calibrated to the roles humans and AI will play when using these skills, as exemplified in the figure. When these basic intellectual skills are solid, the central educational objective becomes to develop engineers who can solve challenging real-world problems through a blend of human creativity and trustworthy use of AI tools. Future engineers must be able to formulate appropriate instructions, interpret AI-generated solutions, and evaluate their correctness and robustness. This includes being the human-in-the-loop who makes ethical tradeoffs, explains system behavior, identifies hidden assumptions, and detects weaknesses arising from poor input data or AI hallucinations. These skills belong to the upper three cognitive levels in Figure 1.

Achieving this objective in the wireless area requires a solid foundation in signals and systems, probability theory, electromagnetics, and communication theory, including radio propagation models and performance metrics. Within each topic, education should be structured around three elements: the underlying engineering problem, theoretical solutions in simplified scenarios that offer interpretability and abstraction, and an explicit discussion of practical regimes in which these solutions may break down or require refinement.

Preprints.org (www.preprints.org) | NOT PEER-REVIEWED | Posted: 16 April 2026 doi:10.20944/preprints202604.1172.v1

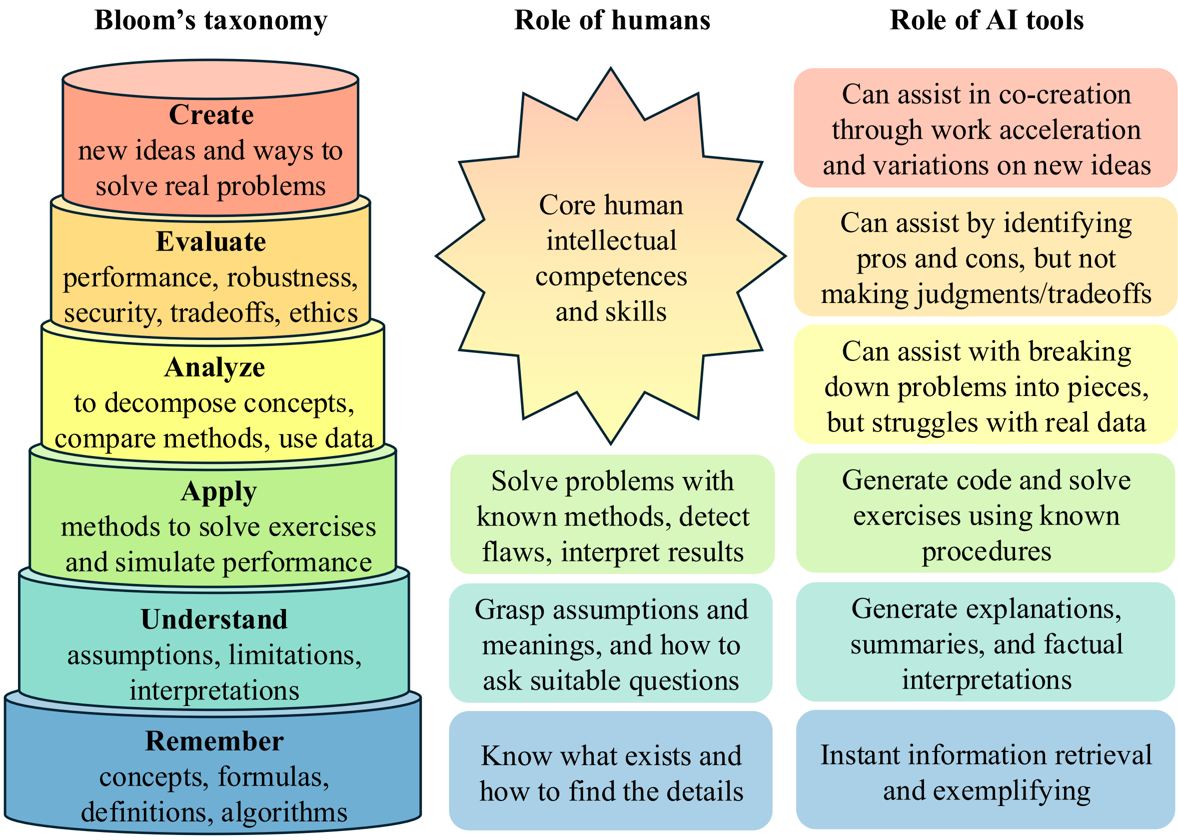

Figure 1. Bloom's classical taxonomy divides educational learning objectives into six cognitive levels. AI tools can automate many lower-level work tasks, so education should focus on elevating students to the highest levels, where human competencies and skills will remain essential in workplaces.

A substantial portion of the foundational material can be studied individually with the support of AI-based tutors that are aware of predefined learning objectives and carefully configured and tested by human teachers. There is growing evidence that students are more inclined to ask questions and explore uncertainties when interacting with non-judgmental AI tutors [6]. Such systems allow students to engage with the material flexibly—across time, location, and modality—while instructors can deploy quizzes and other learning analytics to monitor progress, trace and visualize learning trajectories, and provide personalized guidance [3]. This approach naturally enables personalized work plans and homework assignments, with variation in permitted technological aids.

Maintaining student motivation in a highly digital and distraction-rich environment is a major pedagogical challenge. Providing physical coworking spaces on campuses and fostering cultures in which AI agents manage social media accounts during study hours might help mitigate procrastination. In-person classroom activities remain essential, but not for delivering new theoretical content. The goal is to practice verbalizing, utilizing, challenging, and stress-testing knowledge to elevate both intellectual and practical skills beyond what individualized study can achieve [6]. Seminar-style teaching plays a central role by creating a forum for reflecting on models, assumptions, and AI-generated results, and for internalizing professional norms regarding what is reasonable, robust, and trustworthy in real wireless systems. AI tools can support the instructor in real-time by documenting discussions, injecting alternative perspectives, and flagging inconsistencies in arguments in technical debates.

Collaborative project work further reinforces these goals by engaging students in solving communication-related problems under more realistic conditions than those assumed when teaching the theoretical foundations. This includes working with advanced simulators of practical systems as well as conducting hands-on experiments with software-defined radios, antennas, and over-the-air measurements. AI tools lower the barriers to experimentation by actively guiding students in setting up hardware and managing software interfacing, so students can focus on the actual learning goals. Access to well-equipped makerspaces where students can play with hardware and software under teacher supervision will become a key selling point for studying at universities in the future. Through these activities, students are learning to apply their knowledge beyond toy examples [2] and are exposed to practical impairments such as synchronization errors, external interference, and phase noise, as well as practical design metrics that balance bitrates and latency and incorporate new data-driven semantics. Such projects can be conducted as internships at companies or research groups, and the outcomes form a portfolio to show prospective employers.

These experiences highlight the distinction between the static conceptual building blocks of communication—such as modulation, coding, and multiplexing—and their scenario-dependent implementations. Students gain hands-on experience in characterizing and mitigating impairments by collecting data, training their own AI models, and reflecting on the outcomes [6]. The final assessment and grading should evaluate cognitive skills, human agency, and technical craftsmanship, including the ability to solve open-ended tasks using all available tools and to explain the approach, the solution, and its inherent limitations. Teachers might need support from AI and peer correction to achieve scalability when implementing the examination, while remaining the ultimate decision-makers. While traditional exams ask students to find answers to specific questions, future examinations might also provide AI-generated solutions and ask students to expose their weaknesses: hidden assumptions, errors, and security issues.

By the end of the educational program, graduates should be proficient in using AI tools to solve practical problems in the communications domain, and also experienced in AI-mediated learning of new concepts and skills. The latter is essential for rapid onboarding in workplaces, and for being a lifelong learner who can continue upskilling through AI-customized teaching materials that convey new knowledge and methods in a manner tailored to the individual's current expertise [4].

Preparing students for real-world engineering practice requires more than technical competence. Education must also address the ability to solve complex problems collaboratively, where team members hold distinct responsibilities and discipline expertise, yet depend on shared outcomes. This includes managing social interactions, ensuring robustness and security of jointly developed solutions—potentially supported by adversarial AI methods—and confronting "wicked" problems that are ill-defined and lack a single optimal solution. In such cases, engineers must learn to balance technical performance with considerations of sustainability, regulation, and policy through informed, human-centered judgment. In summary, AI will make the learning experience more rewarding through deeper intellectual stimulation and more experimentation.

AI-Assisted Research

The scientific landscape is experiencing a profound structural shift as research methodologies transition from traditional manual exploration to automated, AI-driven discovery. This transition is characterized by the emergence of open-source ecosystems that integrate reasoning large language models (LLMs) with agentic frameworks capable of executing the end-to-end research lifecycle, as schematically shown in Figure 2. Projects such as the AI Scientist [7] can autonomously generate hypotheses, conduct literature reviews, execute experiments, and synthesize findings into manuscripts with a quality similar to peer-reviewed papers. This transformation is not just an improvement in efficiency, but a fundamental change in the way research is conducted, moving toward a model of "collaborative intelligence" where autonomous agents act as force multipliers for human intellect. Researchers become supervisors of a scalable digital workforce that operates continuously on their behalf. This paradigm shift redefines every phase of the scientific lifecycle, as explored below.

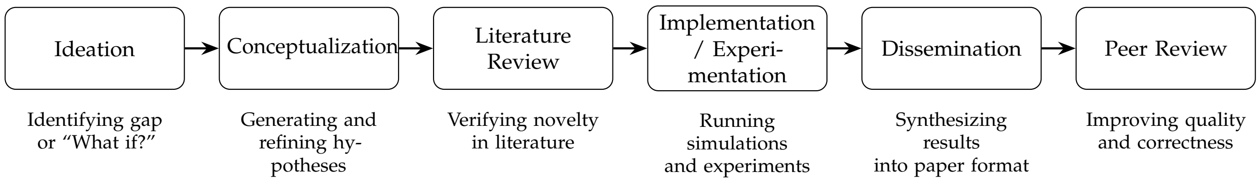

Figure 2. Typical steps in the end-to-end research lifecycle.

During ideation, AI serves as a partner for "combinatorial innovation", rapidly connecting concepts from different fields. However, we must remember that these models are statistical engines trained on past data, which inherently predisposes them to risk aversion. Recent studies confirm this limitation: the authors of [8] caution that, while AI generates novel ideas, they often lack diversity and tend to stick to safe, standard themes. Furthermore, it was found in [9] that while AI is now more creative than the average person, it still trails behind the top 1% of human thinkers. This creates a new bottleneck: since the downstream execution of research can now be largely automated, we face a shortage not of labor, but of high-quality, non-consensus ideas. Thus, our goal should not be to simply use AI to speed up daily work, but to help us find the rare, transformative ideas that actually move the field forward. Silver and Sutton [10] argue that future AI systems that generate their own experience, rather than learning solely from human-generated data, will eventually achieve this.

Automated literature review is currently one of the most mature applications of AI in research, primarily because it solves the massive retrieval and synthesis challenges that humans struggle to scale. Rather than manually searching keywords and sifting through unmanageable lists of citations, research agents can semantically scan thousands of papers in seconds to extract relevant data and synthesize answers to specific inquiries. This capability enables a broader and less biased review process—mitigating our natural tendency to prioritize well-known authors or institutions—while uncovering valuable methodologies from adjacent fields that standard search techniques would likely miss.

The advent of agentic coding systems has dramatically lowered the barrier to entry for implementing and benchmarking new algorithms in simulation environments and testbeds. However, the success of this approach relies on a critical infrastructure: open-source tools. These tools serve as the necessary "symbolic toolbox" that transforms wireless research into a verifiable domain. Much like agentic coding succeeds by fusing statistical generation with logical execution (compilers and shells), AI-assisted wireless research requires open-source stacks to provide an execution environment. This allows agents to run tests, identify hallucinations, and iteratively correct their own mistakes. To fully industrialize this process, we must extend this "toolbox" concept to physical infrastructure. Ideally, entire labs and over-the-air (OTA) testbeds should be exposed as remote, agent-accessible tools, mirroring initiatives like the Science Context Protocol (SCP). By transforming facilities into *agent-ready* assets—prioritizing machine-parsable schemas over human-readable manuals—agents could autonomously orchestrate complex cross-lab campaigns, seamlessly chaining simulations, measurements, and OTA validation in a reproducible manner. In addition, providing an agentic system with *skills*—expert knowledge on how and when to use such assets—will become increasingly important. Regarding the AI capabilities required here, our assessment suggests that frontier models paired with robust retrieval-augmented generation (RAG) pipelines are sufficient for most wireless research use cases; the primary benefit of telco-specific fine-tuning lies not in capability, but in distilling these skills into smaller, faster, and cheaper models.

The use of AI in the dissemination phase belongs to the broader field of technical publication, which will be discussed in the next section.

Technical Publication

AI is infiltrating every aspect of the publication process, from creation to review to consumption. This applies to all kinds of technical publications: journal papers, conference papers, standards, and blog posts.

 doi:10.20944/preprints202604.1172.v1

Near-Term Considerations

Making Content AI-Friendly

We increasingly consume papers with the help of AI to summarize the contributions, explain a derivation, compare results to prior work, or deconvolve complicated language in a standard. This shift in consumption habits has implications for how we should write and package our future publications. Current AI models can read and reason about LaTeX source files, structured text, and well-captioned figures with high fidelity. They struggle, however, with equations embedded as images, as found in many standards documents, and with figures whose captions do not explain what is being shown or what to conclude from the result. Simulation parameters buried in prose are easily missed, and assumptions might even be hidden between the lines. Without access to source code or underlying data, an AI system must guess, rederive, or reproduce using trial-and-error from visual inspection of plots. The solution is to make publications as machine-readable as possible: provide simulation code, share figure data, write self-contained figure captions, and write mathematics as LaTeX-like text rather than as embedded images. Recent work advocates going further by producing research findings in structured, machine-readable formats from the start [11]. Formal proof documents written in languages such as Lean could supplement traditional derivations, enabling AI verification. These practices also benefit human readers by promoting reproducibility and transparency.

AI-Assisted Generation and Editing

AI tools can assist at every stage of manuscript preparation, from organizing scattered notes into a coherent outline to expanding an outline into paragraphs and polishing near-final text for consistency. They can also generate complete text and graphics from a basic set of ideas. AI-generated text, however, tends toward safe, hedging language and formulaic structure that lacks the writer's voice, and thoughtful writing practices developed over a career can erode if we defer too heavily to AI for drafting. We should develop our own writing voice before relying on AI assistance, train the AI for our style, and rewrite AI output rather than accepting it as-is.

Building a manuscript can follow the same agentic workflow as building software. Table 1 maps the phases of an agentic coding pipeline to their paper writing analogs. Each phase can be handled by a specialized agent, and the researcher reviews at phase boundaries, just as a developer approves each stage in an agentic coding tool.

Table 1. Mapping the agentic coding workflow to the paper writing process. Each phase of the coding pipeline has a direct analog in manuscript preparation, and each can be handled by a specialized AI agent with human review at phase boundaries.

Software agent	**Publishing agent**	**Description**
Gather context	Assemble inputs	Collect derivations, simulation results, slides, and code from the research phase
Plan	Outline	Draft paper structure, section organization, key arguments
Implement	Draft	Expand outline into full paragraphs and sections
Test/Verify	Check	Verify proofs, audit notation, enforce style rules
Review	Critique	Generate reviewer feedback; identify weaknesses
Iterate	Revise	Edit manuscript to address review feedback
Commit/Deploy	Submit	Format final manuscript; prepare submission package

Review Processes

AI is a powerful tool for reviewing manuscripts. One approach is to configure an AI system with personal style preferences and writing guidelines, then audit a manuscript for violations. This

 doi:10.20944/preprints202604.1172.v1

kind of enforcement is tedious for humans but straightforward for machines. A second approach is to use AI as "reviewer #2": provide a manuscript and ask for critical feedback, then address the feedback manually. Several startups and free tools now offer automated review, and general-purpose AI tools with well-crafted prompts can achieve useful results [12]. Systems like the automated reviewer module described in [7] have demonstrated that agents can autonomously critique manuscripts and grade them against conference standards with near-human accuracy. Some conferences are already integrating AI into their peer-review process. AAAI 2026 provides authors with an AI-generated review of their submission, while ICLR 2025 used AI to give feedback to reviewers on the quality of their own reviews. Journals might take this further by requiring authors to address a set of AI-generated reviews as a prerequisite for entering peer review. Such a requirement would raise the baseline quality of submissions, reduce the burden on human reviewers, and eliminate the incentive for reviewers to submit a hastily generated AI review.

Longer-Term Directions

End-to-End Publication Generation

The agentic workflow in Table 1 assumes that the research is already complete. Emerging systems go further by automating the research itself, from hypothesis generation through experimentation, and then feeding the results into the publication pipeline. The first version of the AI Scientist project demonstrated this full loop autonomously, and a later version produced the first entirely AI-generated paper passing serious peer review [7]. These systems work best when the problem is well-defined, and the evaluation criteria are unambiguous. In other cases, a human could supervise and steer the process as desired. As they mature, they will reshape expectations about the pace and volume of publication, and may drive the creation of new venues and review processes designed specifically for AI-generated research.

Semantic Compression and Rendering

As AI assumes a larger role in both generating and consuming papers, a fundamental question emerges: What is a paper? If both the writing and the reading are mediated by AI, the canonical form of a publication could shift from polished prose to a compressed semantic representation, a structured encoding of key ideas, assumptions, methods, results, and conclusions. In essence, AI prompts, code, and data. The reader's AI assistant then renders the paper in a format tailored to the reader's preferences and expertise, much as a PDF viewer renders a document from an underlying structured description today. This shift would also resolve many of the challenges discussed earlier, from machine readability to agentic workflows, since the underlying representation is designed for AI from the start.

Enhanced Publishable Units

Regardless of the format of future scientific reports, we must question what level of intellectual content merits peer-reviewed publication, and patents for that matter (as discussed later). If nothing changes, more efficient workflows will lead to vastly higher productivity and a race to the bottom on minimum publishable units. The individual units will then have a negligible impact on other researchers and their AI agents. The research community must eventually come together to create new, stricter definitions of what is considered novel and important, in a format that is readable for AI reviewers and the agents that help researchers to prepare publications, so the rules can be applied consistently within the research field.

Usage of AI in Production-Ready Environments

The use of AI is transitioning from experimental proofs-of-concept to production-grade deployments across the telecommunications value chain: customer engagement, where AI agents reshape support and field engineering; network solutions, spanning external orchestration and embedded networking intelligence; and product development, from AI-assisted coding to future generative hardware design. The following subsections trace this evolution from customer-facing applications through network intelligence to the tools reshaping how products are built.

Customer Engagements

LLMs trained on vast internet corpora have reached human-level performance in text and voice interaction, while real-time video generation is emerging. For end users and the telecommunications ecosystem at large, this capability creates three distinct production fronts.

Consumer Avatars

Operators are beginning to deploy consumer-facing AI avatars, i.e., conversational agents augmented with operator-specific guardrails. This is typically done to deliver personalized assistance without deep integration into backend billing or provisioning systems—just an operator-branded app to lower churn. Because such solutions require relatively shallow system coupling, they represent a low-barrier entry point for operators seeking rapid customer-experience gains and early competitive differentiation.

Customer and Enterprise Support

More deeply integrated are RAG pipelines connected to product catalogs, troubleshooting databases, customer-history repositories, and other internal data sources. Leveraging use-case-specific data, these systems power life-like voice agents in live call centers. An emerging complication is the rise of consumer-side AI agents that can negotiate with the operator's AI system on the customer's behalf. When both sides deploy AI representatives, interactions converge toward Pareto-optimal outcomes far faster than human-mediated exchanges, raising strategic questions about the future of AI agents in customer service. The same methodology extends naturally to business-to-business enterprise account management, quote generation, and contract-negotiation assistance.

Field Engineering Support

For field technicians, AI systems combine general knowledge with operator- and vendor-specific fine-tuning and RAG into technical documentation. Paired with digital-twin platforms that overlay real-time asset data through augmented reality, these tools significantly accelerate troubleshooting and reduce truck rolls. Importantly, they enable less experienced technicians to resolve complex issues, acting as a natural upskilling engine in the commercial world.

Network Solutions

AI-for-Networks

AI-for-networks refers to intelligence that manages, optimizes, and automates network infrastructure from the outside. It typically operates on telemetry, configuration, or policy rather than within the data or control plane itself. In this context, AI-driven network planning and digital twins are maturing rapidly. Tasks such as capacity planning, site selection, and spectrum optimization are increasingly guided by large-scale, agent-based simulations. These tools enable operators to evaluate complex "what-if" scenarios before committing capital. Furthermore, digital-twin representations of live networks effectively close the loop by feeding measured performance data back into planning models, ensuring continuous, data-driven optimization. These capabilities increasingly automate tasks traditionally performed by network planning engineers, suggesting significant workforce implications.

Beyond planning, operations are shifting from closed self-organizing networks toward open, multi-vendor programmable architectures. Frameworks like the O-RAN Alliance's Service Management and Orchestration (SMO) have reached production, with major Tier-1 operators successfully deploying third-party rApps on commercial networks. Current capabilities focus on anomaly detection, key performance indicator (KPI) monitoring, energy savings, and light closed-loop automation. Independent validations of these autonomous operations—such as recent Level 4 Radio Access Network (RAN) energy efficiency certifications—underscore a growing commitment to standardized, AI-supported network management.

Finally, AI is transforming network operations centers (NOCs) by automating root-cause analysis, predictive maintenance, and self-healing. Benchmarks indicate that smaller, domain-tuned models often outperform large frontier models on specialized telecom tasks [13]. However, intent-to-configuration translation remains a critical gap. Since current models struggle to output schema-compliant configurations, fully autonomous closed-loop management remains an engineering challenge, albeit without fundamental barriers.

AI-in-Networks

AI-in-networks refers to embedded intelligence within network functions, operating directly on signal processing, resource allocation, and mobility management. Within the RAN, the 3rd Generation Partnership Project (3GPP) has progressively specified AI use cases across successive releases. Following the foundational study in Release 18, which prioritized channel state information (CSI) compression, beam management, and positioning, Release 19 advanced these capabilities into normative specifications. Building on this momentum, Release 20 was a milestone with the approval of the "NR air interface Phase 2" study item in mid-2025. This phase transitions the industry toward two-sided model deployments for tasks like CSI compression, directly addressing complex operational challenges, including model pairing, fallback mechanisms, version synchronization, and inter-vendor cooperation.

Despite ongoing standardization, translating theoretical gains into real-world performance is challenging, as simulated improvements rarely generalize across heterogeneous global conditions. Consequently, the pervasive deployment of native AI in the RAN will likely take over a decade. This trajectory mirrors the slow evolution of autonomous driving from modular, rule-based stacks to unified, end-to-end neural engines.

Beyond the radio edge, intelligence is advancing in the 5G Core and end-to-end orchestration. Capabilities like AI-supported traffic steering, slice management, and policy optimization are nearing production. However, deployments remain largely vendor-proprietary, with core standardization lagging behind the RAN. Cross-domain frameworks, such as the TM Forum's Autonomous Networks model, target Level 4 autonomy by 2030. Although recent certifications show momentum, most operators remain at Levels 2–3. As a result, current industry consensus favors high-value, scenario-specific autonomy rather than attempting end-to-end automation.

Product Development and Innovation

Intellectual Property

The U.S. Patent and Trademark Office issued revised guidance in November 2025 clarifying that AI-assisted inventions are patentable under standard legal criteria. AI systems are treated as tools analogous to laboratory equipment, as long as a natural person has conceived the invention [14]. This framing sidesteps deeper questions: if inventions emerge from iterative prompting of stochastic models, the boundary between human conception and machine generation becomes increasingly blurred——a tension that patent law has yet to fully resolve. For standards authoring, AI can generate OpenAPI specifications and code snippets, but full-standard authoring remains constrained by non-machine-parsable formats (e.g., equations embedded as images in current 5G specifications). Once AI can draft standards in their entirety, the future of standards should be revisited, but we are still a few years away from that.

Code Generation

Modern coding tools, such as Claude Code and Cursor, have moved beyond autocomplete into fully agentic frameworks that operate with increasing independence across planning, implementation, and testing stages over long periods of time. Porting of significant code bases using AI-assisted tools has demonstrated order-of-magnitude speedups relative to manual estimates, illustrating the

Preprints.org (www.preprints.org) | NOT PEER-REVIEWED | Posted: 16 April 2026 doi:10.20944/preprints202604.1172.v1

productivity gains achievable on well-defined and verifiable tasks. We expect that huge parts of 6G implementation will be done with the help of such tools, probably being the most significant application of AI in the field of telecommunications in the next few years.

Toward Self-Synthesizing Networks

The most ambitious frontier is AI that generates novel hardware, protocols, and network architectures—not for research but for deployments. The main bottleneck is data scarcity: telecommunications hardware design data is proprietary and insufficient for training generative models at scale. First evidence, however, emerges that generative AI can be used for radio frequency integrated circuit (RFIC) designs, allowing for the discovery of non-intuitive circuit architectures that outperform human designs while compressing design cycles from months to weeks [15]. Foundational telecom-specific generative models remain in early stages, and the path from isolated design breakthroughs to systematically self-synthesizing networks is long.

Final Thoughts

While the integration of AI into every facet of research seems inevitable, we must recognize that throwing AI at every process is not inherently beneficial. There remains, for now, a profound, irreplaceable value in structured, unplugged cognitive work and in understanding fundamental, unbeatable performance limits. In highly abstract or emerging domains—such as quantum communications or advanced multi-linear algebra—current AI tools often provide surprisingly little assistance. Our practical experience in these nascent fields shows that some of the most profound insights still emerge from human researchers working exclusively with pen and paper, deliberately disconnected from digital distractions. This highlights a critical challenge for the future: the self-preservation of human cognitive ability. As we embrace lifelong learning, we must actively design educational frameworks and professional habits that force both students and seasoned engineers to practice deep, unaided thinking. If we rely entirely on AI for the heavy lifting, we risk atrophying the very skills required to comprehend and verify the systems we are building.

Furthermore, we must critically evaluate what we actually want AI to achieve in the research process. The ultimate goal of science is not to marginally improve the quantity and quality of average papers, but to generate the rare, outlier breakthroughs that redefine our understanding and practices. Currently, AI models excel at interpolation within the boundaries of existing knowledge. Even seemingly novel systems, such as LLM-guided genetic algorithms, often resemble sophisticated recombination rather than true, paradigm-shifting creativity—though one might argue this mirrors human ideation to some extent. The grand open challenge is how we can train and steer these systems to move beyond mere interpolation to genuine discovery. Agentic systems that learn from their own experience with the world [10] might be able to achieve this.

As AI dramatically shrinks the time and cost required to explore new research directions, the fundamental bottleneck in innovation is shifting. AI will execute relentlessly in whatever direction it is pointed, entirely indifferent to whether a problem is actually worthwhile, impactful, or meaningful. Consequently, the human researcher's primary value will shift from solving problems to identifying the right problems to solve and to balancing technical, economic, legal, and ethical constraints. Ideally, the mechanical execution performed by AI will free up our cognitive bandwidth to spend more time deeply considering the overarching challenges that truly matter. However, this shift introduces new risks of systemic inequality. Just as academia currently struggles with disparities in computational hardware, the future research landscape may be sharply divided by access to it. The ability to produce top-tier research and breakthrough papers may become heavily gated by who can use the most advanced frontier AI models.

Ultimately, while the transition period will likely be turbulent, we maintain a cautious optimism about the future of our profession. The near term will undeniably be marked by significant friction—including challenges to research integrity, narrowing of scientific scope, and painful disruption to entry-level employment as routine engineering tasks are automated. Yet, if we navigate this disruption

 doi:10.20944/preprints202604.1172.v1

thoughtfully, the long-term impact on the wireless communications ecosystem can be mostly positive. By stripping away the rote mechanics of our daily work, AI will ultimately elevate the human role, challenging us to become better visionaries, deeper thinkers, and the true architects of future connectivity.

Acknowledgments: E. Björnson thanks his research group and the SweWIN center, funded by Vinnova (2023-00572). J. Hoydis thanks his colleagues F. Aït Aoudia, S. Cammerer, L. Maggi, G. Marti, and A. Keller. R. Heath thanks the support from the National Science Foundation under Grant No. NSF ECCS-2414678 and the Army Research Office under Grant W911NF2410107.

References

1. Saad-Falcon, J.; Narayan, A.; Akengin, H.O.; Griffin, J.W.; Shandilya, H.; Lafuente, A.G.; Goel, M.; Joseph, R.; Natarajan, S.; Guha, E.K.; et al. Intelligence per Watt: Measuring Intelligence Efficiency of Local AI. *arXiv preprint arXiv:2511.07885* **2025**.
2. Rak, G. How to Stay Ahead of AI as an Early-Career Engineer. *IEEE Spectr.* **2025**.
3. Viberg, O.; Poquet, O.; Kovanovic, V.; Khosravi, H. Fostering Human Agency in Age of AI: A Learning Analytics Perspective. *J. Learn. Analytics* **2025**, *12*, 1–7.
4. Bekmanova, G.; Ongarbayev, Y.; Somzhurek, B.; Mukatayev, N. Personalized Training Model for Organizing Blended and Lifelong Distance Learning Courses and Its Effectiveness in Higher Education. *J. Comput. Higher Educ.* **2021**, *33*, 668–683.
5. Clark, L.A. Reframing Bloom's for the Age of AI: A White Paper for Future-Ready Educators. Anthology, 2025.
6. Cabrera, A. What and How to Teach When Google Knows Everything and ChatGPT Explains It All Very Well. Georgia Institute of Technology, Office of the President Blog, 2025.
7. Yamada, Y.; Lange, R.T.; Lu, C.; Hu, S.; Lu, C.; Foerster, J.; Clune, J.; Ha, D. The AI Scientist-v2: Workshop-Level Automated Scientific Discovery via Agentic Tree Search. *arXiv preprint arXiv:2504.08066* **2025**.
8. Si, C.; Yang, D.; Hashimoto, T. Can LLMs Generate Novel Research Ideas? A Large-Scale Human Study with 100+ NLP Researchers. In Proceedings of the Proc. Int. Conf. Learn. Representations (ICLR), 2025.
9. Bellemare-Pépin, A.; Lespinasse, F.; Thölke, P.; Harel, Y.; Mathewson, K.W.; Olson, J.A.; Bengio, Y.; Jerbi, K. Divergent Creativity in Humans and Large Language Models. *Sci. Rep.* **2026**, *16*, 1–18.
10. Silver, D.; Sutton, R.S. Welcome to the Era of Experience, 2025. Preprint of a chapter to appear in *Designing an Intelligence*, edited by George Konidaris, MIT Press (forthcoming).
11. Kuhn, T.; et al. Rethinking the Production and Publication of Machine-Readable Expressions of Research Findings. *Sci. Data* **2025**, *12*, 1–17.
12. Lin, J.; et al. Large Language Models for Automated Scholarly Paper Review: A Survey. *ACM Inf. Fusion* **2025**, *124*, 1–17.
13. GSMA. GSMA Open-Telco LLM Benchmarks 2.0: The First Dedicated LLM Evaluation for Telecoms. Hugging Face Blog, 2025.
14. U.S. Patent and Trademark Office. Revised Inventorship Guidance for AI-Assisted Inventions. Federal Register, 2025. vol. 90, no. 228.
15. Karahan, E.A.; Liu, Z.; Gupta, A.; Shao, Z.; Zhou, J.; Khankhoje, U.K.; Sengupta, K. Deep-Learning Enabled Generalized Inverse Design of Multi-Port Radio-Frequency and Sub-Terahertz Passives and Integrated Circuits. *Nat. Commun.* **2024**, *15*, 1–13.

Short Biography of Authors

Emil Björnson[F'20](emilbjo@kth.se) is a Professor of Wireless Communication at the KTH Royal Institute of Technology, Sweden, where he also leads the Department of Communication Systems and the SweWIN research center. He received his master's degree from Lund University, Sweden, and his Ph.D. from KTH. He has co-authored four textbooks, published 400 peer-reviewed papers, and created 70 open-source software repositories. He is a Clarivate Highly Cited Researcher and co-host of the YouTube channel "Wireless Future".

Preprints.org (www.preprints.org) | NOT PEER-REVIEWED | Posted: 16 April 2026 doi:10.20944/preprints202604.1172.v1

Mischa Dohler [F'14](mischa.dohler@ericsson.com) is the Vice President of Emerging Technologies at Ericsson Inc. in Silicon Valley. He received his diploma degree in electrical engineering from TU Dresden, Germany, and his Ph.D. in telecommunications from King's College London, U.K. He was previously a Chair Professor of Wireless Communications at King's College London, where he directed the Centre for Telecommunications Research. He is a Fellow of the Royal Academy of Engineering and a Fellow of the Institution of Engineering and Technology.

Jakob Hoydis[F'24](jhoydis@nvidia.com) is a Distinguished Research Scientist at NVIDIA working at the intersection of machine learning and wireless communications. He received his diploma degree from RWTH Aachen University, Germany, and his Ph.D. from Supélec, France. He co-authored "Massive MIMO Networks: Spectral, Energy, and Hardware Efficiency" and is one of the core developers of Sionna, a GPU-accelerated open-source software library for next-generation communication systems.

Robert W. Heath Jr.[F'11](rwheathjr@ucsd.edu) is the Charles Lee Powell Chair in Wireless Communications at the University of California, San Diego. He has co-authored "Foundations of MIMO Communications" and authored "Introduction to Wireless Digital Communications: A Signal Processing Perspective." He is an elected member of the U. S. National Academy of Engineering.