



6G

Next G Alliance Report:
**Sustainable AI in Telecom:
Promises and Challenges in 6G**

As a leading technology and solutions development organization, the Alliance for Telecommunications Industry Solutions (ATIS) brings together the top global Information Communications Technology (ICT) companies to advance the industry's business priorities. ATIS' 150 member companies are currently working to address network reliability, 5G, robocall mitigation, smart cities, artificial intelligence (AI)-enabled networks, distributed ledger/blockchain technology, cybersecurity, IoT, emergency services, quality of service, billing support, operations and much more. These priorities follow a fast-track development lifecycle from design and innovation through standards, specifications, requirements, business use cases, software toolkits, open-source solutions, and interoperability testing.

ATIS is accredited by the American National Standards Institute (ANSI). ATIS is the North American Organizational Partner for the 3rd Generation Partnership Project (3GPP), a member of the International Telecommunication Union (ITU), as well as a member of the Inter-American Telecommunication Commission (CITEL).

For more information, visit www.atis.org. Follow ATIS on [Twitter](#) and on [LinkedIn](#).

The ATIS 'Next G Alliance' is an initiative to advance North American wireless technology leadership over the next decade through private-sector-led efforts. With a strong emphasis on technology commercialization, the work will encompass the full lifecycle of research and development, manufacturing, standardization, and market readiness.

As the global climate crisis intensifies, governments, industries, and communities worldwide face an increasing pressure to accelerate efforts toward achieving Net Zero emissions. Although the rapid adoption of Artificial Intelligence (AI) is transforming the Information and Communications Technology (ICT) landscape, it presents both opportunities and challenges. While AI offers powerful tools to optimize energy use, enhance supply chain efficiency, and accelerate renewable energy adoption, its computational demands—particularly during model training and deployment—pose significant environmental challenges, such as increased energy consumption and carbon emissions. This paper examines the dual impact of AI on Net Zero goals by exploring two key concepts: “AI for Sustainability” and “Sustainable AI.”

To mitigate AI’s environmental impact across the ICT value chain, this paper underscores the importance of evaluating the entire lifecycle of AI systems, from inception to retirement. It outlines strategies to reduce embodied emissions in AI infrastructure, transition to renewable energy sources, and minimize the energy consumption of Machine Learning (ML) processes. Additionally, it highlights how AI can enhance energy efficiency across key areas, including Radio Access Networks (RAN), core networks, User Equipment (UE), and data centers. The critical role of Key Performance Indicators (KPIs) and Key Value Indicators (KVI) in assessing the effectiveness of these strategies is also emphasized.

AI-driven techniques are emerging as essential tools to optimize energy and sustainability performance in the rapidly evolving telecommunications sector. This paper explores forward-looking aspects such as the integration of AI with quantum computing, Non-Terrestrial Networks (NTNs), and opportunities for AI to enhance spectral efficiency and enable intent-based automation in networks.

Summarizing the challenges of managing energy consumption, carbon emissions, water usage, and electronic waste, this paper provides actionable recommendations for leveraging AI to analyze radio network conditions and user preferences to enhance sustainability performance while maintaining Quality of Service (QoS). While AI holds immense potential to optimize the environmental sustainability of next-generation networks, the paper underscores the critical need to address its broader social and environmental impacts. We invite you to explore this comprehensive analysis and discover how thoughtful implementation of AI can drive meaningful progress toward a sustainable future.

TABLE
OF
CONTENTS

	Foreword	i
	Executive Summary	ii
1	Introduction	1
1.1	Understanding Sustainable AI and AI for Sustainability	2
1.2	The Role of AI and its Environmental Impacts	2
1.3	Understanding the Energy Consumption of AI	3
1.3.1	Environmental Impacts of AI's Operational Stage	3
1.3.2	Strategies for Powering AI to Reduce Environmental Impacts	3
1.4	Sustainability in Telecom Enabled by AI	4
2	Sustainable AI in Telco Domains	6
2.1	Radio Access Network	6
2.2	Core Network	7
2.2.1	NWDAF Energy Saving Actions:	7
2.3	User Equipment	8
2.4	Data Center	8
2.4.1	Existing Data Center Sustainability Metrics	9
3	Measuring the Environmental Sustainability of AI	10
3.1	Measuring and Monitoring of AI/ML Energy Efficiency	11
3.1.1	Energy Efficiency of AI/ML Entities	11
4	Techniques for Reducing Energy Consumption in Machine Learning	14
4.1	Model Compression Techniques	14
4.1.1	Pruning	14
4.1.2	Quantization	14
4.1.3	Knowledge Distillation (KD)	16
4.2	Additional Techniques	16
4.2.1	Efficient Architectures	16
4.2.2	Gradient Accumulation	16
4.2.3	Sparse Representations	16
4.2.4	Adaptive Computation	16
4.2.5	Edge Computing	16
4.2.6	Federated Learning	16
4.2.7	Efficient Hardware Utilization	16
4.2.8	Energy-Aware Scheduling	16
4.2.9	Emerging AI Technology	16
4.3	Energy Optimization Techniques for Reducing AI/ML Training and Inference Costs	16
4.4	The Importance of Key Performance Indicators (KPIs) and Key Value Indicators (KVI)	17
4.5	Trade-off Between Energy Consumed Versus Energy Saved	17
5	Future Looking Aspects	18
5.1	AI/ML and the Physical Layer	18
5.2	AI Intent-Driven Techniques in 6G	18
5.3	Leveraging Quantum AI for More Sustainable and Resilient Mobile Networks	18
5.4	AI Sustainability for Non-Terrestrial Networks	19
6	Summary, Challenges, & Recommendations	20
6.1	Challenges	20
6.1.1	Emissions	20
6.1.2	Water Usage	20
6.1.3	Waste Management	20
6.2	Recommendations	20
7	Acronyms & Abbreviations	22
8	References	24
	Acknowledgments	27
	Next G Alliance Reports	30
	Copyright and Disclaimer	31

1 INTRODUCTION

Artificial Intelligence (AI) has rapidly emerged as a transformative force across industries, offering significant potential to enhance efficiencies and drive innovation. In the telecommunications sector, AI is poised to revolutionize network management, optimize service delivery, and create new business opportunities. As highlighted by Omdia, nearly two-thirds of network traffic will likely involve AI by 2030, underscoring its critical role in shaping the future of telecom [1]. However, the challenge lies in balancing the powerful capabilities of AI with the imperative to achieve environmental sustainability goals. While AI can drive tremendous efficiencies and innovations, it also demands substantial computational resources, leading to increased energy, water and resource consumption. This paradox creates a pressing need to find the right equilibrium: harnessing AI to achieve greater efficiencies while mitigating its own environmental impacts.

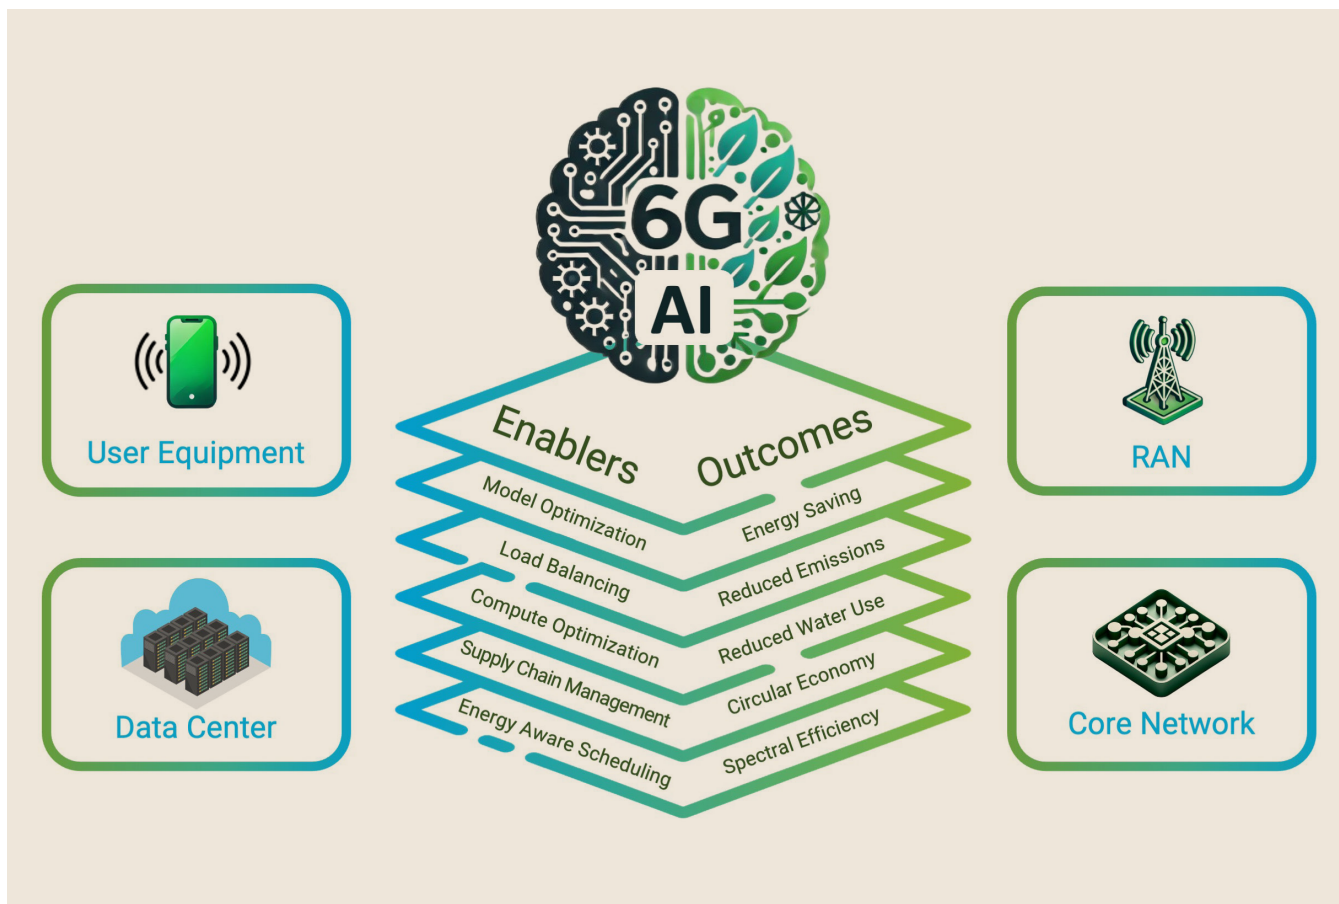


Figure 1: Sustainable AI in Telco: Key Enablers and Outcomes

To address this challenge, companies must adopt a multifaceted approach that aligns AI advancements with their Net Zero goals. This includes developing energy-efficient AI algorithms that reduce the computational load, investing in green data centers that utilize renewable energy sources, and implementing comprehensive lifecycle management practices to minimize the environmental footprint of AI technologies. By integrating these strategies, organizations can leverage AI to achieve their business objectives while remaining committed to sustainability and reducing their overall energy consumption.

The objective of this white paper is to introduce the key concepts of sustainable AI in telecom networks and discuss how AI enabled techniques can optimize environmental sustainability of networks. This includes exploring the correlation between “Sustainable AI” and “AI for Sustainability” as well as discussing the current efforts in the state-of-the-art 5G and the next-generation networks within the Information and Communications Technology (ICT) industry. This white paper does not delve into the societal and economic impacts of AI including those related to responsible AI.

1.1 Understanding Sustainable AI and AI for Sustainability

As organizations seek to reduce their carbon footprint and improve operational efficiency, two concepts are gaining prominence: “Sustainable AI” and “AI for Sustainability.” These terms, while interconnected, represent different but complementary aspects of how AI is driving more sustainable and efficient network operations.

AI for sustainability refers to the use of AI technologies to create more energy-efficient networks while maintaining or enhancing performance. This approach optimizes network operations, reduces energy consumption, and improves resource utilization, all of which contributes to minimizing the environmental footprint of telecommunications operations.

Sustainable AI focuses on the responsible implementation of AI within networks and company operations, ensuring that the technology itself adheres to sustainable practices. In contrast, AI for sustainability utilizes AI to achieve sustainable outcomes, advancing the telecommunications industry toward a greener and more efficient future.

By analyzing AI’s impact on energy consumption, operational efficiency, and resource management, we can identify both its positive contributions and the potential challenges it poses in promoting sustainability in telecommunication networks.

1.2 The Role of AI and its Environmental Impacts

AI is driving technological advancements across various sectors, but its increase in environmental footprint is concerning. The computational power required for AI systems often necessitates large data centers, which consume significant amounts of electricity, primarily sourced from fossil fuels. Research from the University of Massachusetts Amherst revealed that training a single AI model can produce carbon emissions equivalent to five cars over their lifetimes [2]. As AI continues to evolve, balancing technological growth with environmental sustainability becomes an urgent challenge. This predicament emphasizes the need for a more holistic understanding of AI’s lifecycle impacts and the trade-offs between innovation and ecological responsibility.

Despite these concerns, AI has the potential to generate positive environmental impacts by optimizing energy use across different industries. For instance, AI algorithms can enhance the efficiency of renewable energy sources by accurately predicting supply and demand, which minimizes waste and improves grid stability. In telecommunications, AI can help reduce energy consumption through optimized network management, improving the operation of base stations—among the largest energy consumers in cellular networks. Techniques like dynamic resource allocation and predictive maintenance can lead to more efficient resource use and reduced operational costs, showcasing AI’s critical role in creating greener wireless communication systems.

To mitigate the environmental impact of AI, it is essential to assess the entire life cycle of AI systems, from inception to retirement [3]. Figure 2 illustrates the main AI life cycle stages [4] and their relation to the environmental impact assessment, the Life Cycle Assessment (LCA) method [5],[6]. The typical energy-hungry processes of training and inference are separately indicated in Figure 2. Each of the AI life cycle stages can be further separated into substages – e.g., training data preparation and algorithm selection during design and development.

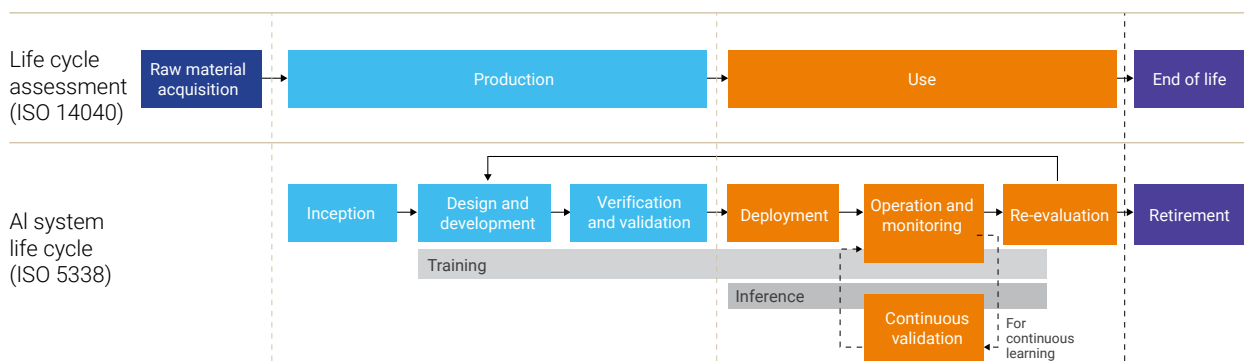


Figure 2: AI System Environmental Impact Assessment [3]

By assessing the environmental impact from the full AI life cycle, specific hot spots can be identified for potential environmental impact reduction. These hot spots are highly dependent on the AI system’s characteristics, its use case, operating location, etc. Energy consumption is one of the main impacting parameters in energy-intensive AI systems. Depending on the system and use case, either training or inference can consume the highest amount of energy. Current estimates indicate that the operational phase, especially the inference stage, is responsible for approximately 80% of its environmental footprint. Addressing this phase’s energy demand requires a coordinated strategy that incorporates efficient data center operations and effective resource allocation to minimize the overall carbon emissions associated with AI technologies.

In addition to climate change, LCA can also provide impact for other environmental impact categories, such as resource use of minerals (e.g., metals and fossils), water use, and land use during the full life cycle of an AI system. When selecting suitable environmental impact categories, the environmental impacts to air, biodiversity, and/or geodiversity can be assessed in a comprehensive and transparent manner.

1.3 Understanding the Energy Consumption of AI

The power required for sustaining the rise of AI is doubling every 100 days [7]. As adoption of AI accelerates, technology companies are reporting up to a 48% increase in greenhouse gas emissions during 2023, driven by construction and operation of data centers supporting AI workloads. As these companies continue to further integrate AI within their products and services, it is going to become increasingly challenging to simultaneously reduce emissions and achieve NetZero emission targets [8]. This challenge underscores the importance of access to power from carbon-free renewable and nuclear sources for the sustainable deployment of AI.

However, in the United States only 40% of the electricity was generated from these sources in 2023 [9]. This challenge is further compounded by the delays faced by developers in securing approvals to connect clean energy projects to the power grid and start delivering the much-needed carbon-free power. As of June 2024, more than 2 terawatts (TW) of clean energy projects are stuck in the permit approval stage [10]. The combined impact of the additional infrastructure and power needed for AI adoption is challenging the technology sector’s ability to achieve Net Zero goals.

1.3.1 Environmental Impacts of AI’s Operational Stage

Training, inference, and retraining of Large Language Models (LLMs) are energy-intensive processes [11]. As the adoption of advanced AI capabilities, such as image, audio, and video generation, grows, the associated energy consumption is also expected to increase. This trend is projected to affect network traffic, according to Omdia’s AI network traffic forecast. By 2030, nearly two-thirds of the network traffic will involve AI content generation, curation, and/or processing, which is likely to drive an increase in mobile network power consumption [12].

Current estimates indicate that approximately 80% of the operational environmental footprint of AI is attributable to the inference phase, with the training phase accounting for the remaining 20%. However, as AI adoption expands across various sectors, the demand for inference—and consequently its associated environmental impacts—are expected to rise as well.

1.3.2 Strategies for Powering AI to Reduce Environmental Impacts

Managing where and when AI workloads are performed can significantly affect the environmental impacts of AI.

AI workloads which can be completed within shorter time periods can be scheduled to align with off-peak hours of the day. During these hours, there is less demand in the utility grid, which can help reduce energy costs, emissions and support the stability of the power grids.

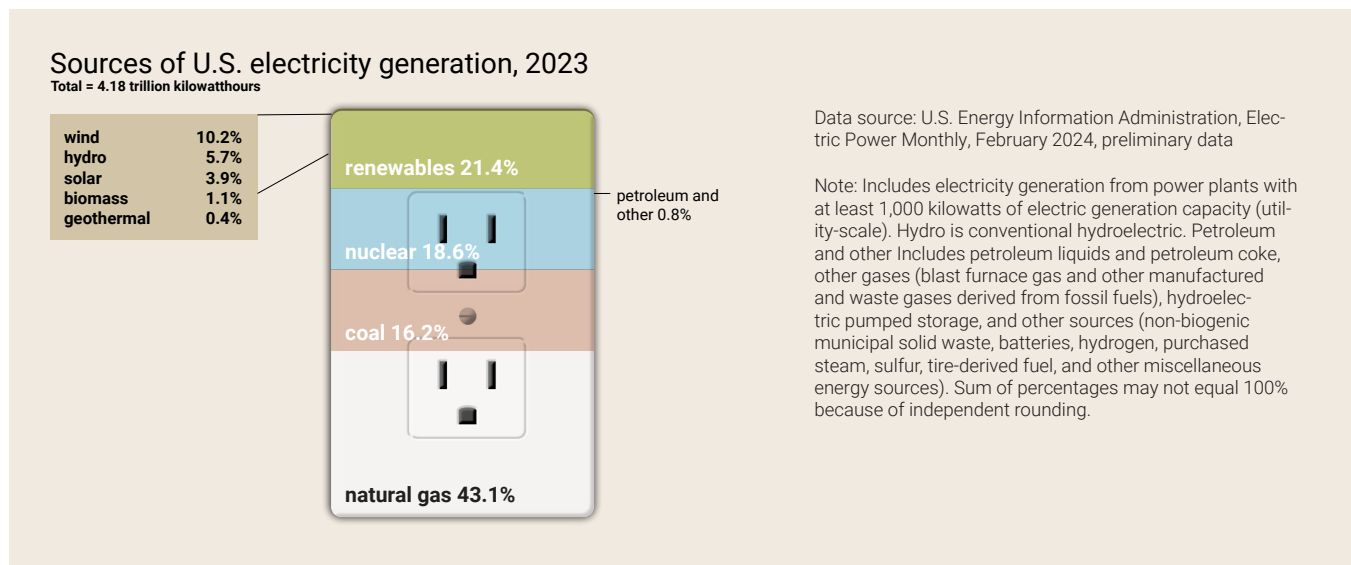


Figure 3: Sources of U.S. Electricity Generation, 2023 [9]




Figure 4: When demand increases, energy costs rise, and fossil fuel plants need to switch on [13]

Based on the data provided by the California Independent System Operator, which manages the California electricity grid, emissions from power generation during winter months are lower than in summer months. Scheduling larger AI projects for execution during the winter months can help reduce greenhouse gas emissions.




Figure 5: Historical CO₂ Trend: Yearly Snapshot of CO₂ Emissions by Month [14]

In 2023, fifteen U.S. states accounted for an estimated 80% of the national data center load. The top five states that host data centers, ranked in descending order, are Virginia, Texas, California, Illinois, and Oregon. Several of these states have a lower CO₂ emissions rate in comparison to the national average. Hosting AI workloads in data centers based on the emission rate of available power can help reduce emissions and enable the sustainable deployment of AI.

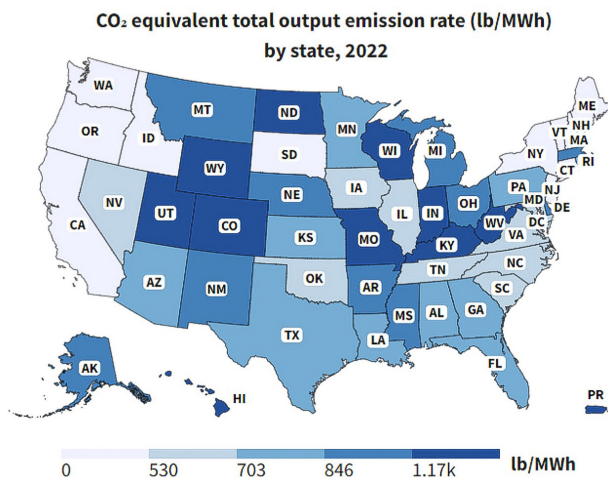


Figure 6: US EPA eGRID Data Explorer output emissions rate by US States [15]

1.4 Sustainability in Telecom Enabled by AI

As many nations have set ambitious targets to reach Net Zero emissions, the telecommunications industry is under growing pressure to enhance the energy efficiency of its networks. In particular, the Radio Access Network (RAN) which is a major power consumer, accounting for more than 80% [16] of the total energy use in mobile networks, whereas the core network accounts for approximately 12%. As network infrastructure continues to transition from 4G to 5G, data center energy consumption has increased, with the increasing movement of RAN components and network functions to data centers in next-generation networks, this trend is expected to accelerate. In this context, AI has emerged as a powerful tool to optimize power usage, minimize operational inefficiencies, and ensure sustainability across next-generation networks.

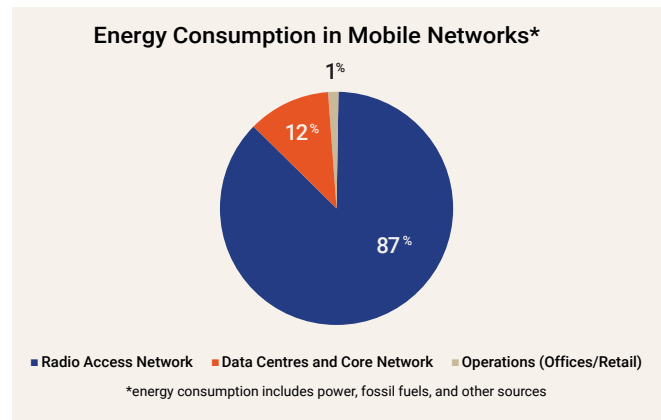


Figure 7: Energy Consumption in Mobile Networks [16]

The United States has set a target to achieve 100% carbon pollution-free electricity by 2035 [17], focusing on the expansion of renewable energy sources like wind and solar to reduce emissions. Similarly, the ICT sector can leverage AI to support its own sustainability transition by optimizing power consumption across both centralized and distributed network architectures. AI's capacity to process vast amounts of data and dynamically adjust to changing conditions can be used to fine-tune network operations, predict traffic loads, and balance resource allocation in real time.

Furthermore, AI can be used to orchestrate the network to balance energy demands by prioritizing power consumption from renewable sources whenever available and minimizing dependence on carbon-based energy. By leveraging predictive

algorithms and real-time analytics, AI can forecast network demand and dynamically allocate workloads to data centers or network nodes that are powered by renewable energy sources. This intelligent energy management ensures that carbon-fueled energy is only utilized during peak demands or when renewable sources are unavailable, ultimately reducing greenhouse gas emissions and aligning the network's energy consumption with broader sustainability goals.

AI is playing an important role in optimizing energy consumption within telecommunications networks. By utilizing AI and ML, telecom network operators can predict traffic patterns and adjust energy usage accordingly, leading to substantial efficiency improvements. For example, Telefónica has implemented over 1,440 energy efficiency projects, including Power Saving Features (PSFs) that deactivate network components during low traffic periods. This initiative, known as "Green Radio," saved up to 26% of energy during low traffic hours through the use of Radio Deep Sleep Mode [18]. Similarly, FarEasTone is utilizing AI-driven solutions to analyze network data, reducing both energy consumption and carbon emissions while maintaining high service standards [19].

Operational Expenses (OPEX) in network operations, largely driven by energy consumption, pose another significant challenge for network operators. Energy costs can account for 20-40% of total network OPEX, with a significant portion attributed to powering passive infrastructure components like batteries and climate control units. To address this, integrated AI and data analytics capabilities, are allowing network operators to monitor passive equipment and optimize energy usage [20]. By automating efficiency scenarios and generating actionable work orders, this system reduces OPEX and CO₂ emissions while enhancing network performance. Additionally, AI is increasingly used in sustainability efforts within data centers, where technologies like air cooling and liquid immersion cooling offer promising solutions for reducing environmental impacts.

2 SUSTAINABLE AI IN TELCO DOMAINS

The integration of AI across four key telco domains – RAN, Core Network, UE, and Data Centers – presents a transformative opportunity to enhance operational efficiency while driving sustainability. AI's ability to analyze vast amounts of data in real time allows for intelligent decision-making that can optimize resource allocation, predict maintenance needs, and enhance customer experiences. Moreover, by embedding AI in these domains, network operators can build more efficient, reliable, and environmentally sustainable networks that significantly reduce energy consumption, lower operational costs, and minimize the environmental impact to align with global sustainability initiatives.

It is important to remember that while incorporating AI in telecom may reduce the environmental impact and operational expenditures, the implementation and maintenance of AI/ML algorithms come at their own cost. Therefore, it is crucial to carefully select ML models that can achieve the required goal with the least amount of inherent overhead. The key consideration is to balance the requirements for model performance and energy consumption of the computing resources. Some of the examples are selecting energy-efficient ML models requiring less parameters to train, leveraging transfer learning instead of training models from scratch, and employing statistical information or data sampling to minimize data storage needed to represent the population data effectively.

which occur during a product's use phase, such as energy consumption. According to the GHG Protocol, embodied emissions are categorized under "Scope 3" emissions, covering indirect emissions that occur in a company's value chain, beyond direct energy use.

In the context of AI systems, embodied emissions are specifically related to the computing infrastructure required for training and deploying AI models. These systems, composed of servers, GPUs, and other hardware, contribute significantly to the overall carbon footprint due to the carbon-intensive processes involved in their production. To account for these emissions over the useful life of the equipment, the total embodied carbon footprint is spread across its usage period. For example, in the case of the BigScience Large Open-science Open-access Multilingual Language Model (BLOOM), embodied emissions accounted for 22.2% of the total carbon footprint, while operational emissions from dynamic and idle energy use contributed 48.9% and 28.9%, respectively [22].

2.1 Radio Access Network

RAN is one of the most energy-intensive components of a mobile network due to the large number of base stations required to provide coverage. Base stations typically operate continuously, consuming significant amounts of

power. Implementing sustainable AI in the RAN is essential because it can drastically reduce energy consumption through intelligent resource management.

In RAN, AI can play a pivotal role in dynamic spectrum management, interference reduction, and energy-efficient network planning. By employing ML algorithms to analyze traffic patterns and predict demand, the network can dynamically optimize spectrum and other network resource allocation to reduce energy consumption during low-usage periods according to real-time traffic patterns. Furthermore, AI can optimize signal processing and interference management, enhancing overall network performance while conserving energy, provided that

the energy consumed by the AI-capable hardware is lower than the energy savings realized by AI-enabled optimization. Implementing these measures can not only reduce the operator's carbon footprint but also lower operational costs, contributing to a more sustainable and economically viable network.

Total carbon footprint of AI includes embodied and operational carbon emissions

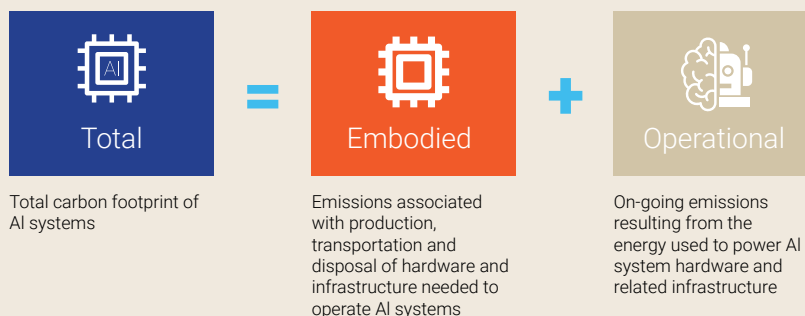


Figure 8: Total carbon footprint of AI [21]

Embodied carbon emissions refer to the total Greenhouse Gases (GHG) emitted during the entire lifecycle of a product's materials and manufacturing processes, prior to its use. This includes emissions from raw material extraction, transportation, manufacturing, assembly, and installation. These emissions are distinct from operational emissions,

2.2 Core Network

The Core Network is the central part of a telecom infrastructure, which handles data routing, network management, and connectivity services. It requires constant and reliable operation leading to high energy usage. Within this domain, AI is leveraged in optimizing data routing, fault management, and load balancing. Predictive analytics can foresee potential network congestion and re-route traffic to less busy pathways, thereby ensuring efficient usage of network resources and minimizing energy waste. Additionally, AI-driven predictive maintenance can detect potential hardware failures and initiate proactive measures, thereby minimizing downtime and energy costs associated with unplanned repairs. By optimizing these processes, operators can reduce the frequency of energy-intensive emergency interventions and maintain stable network performance. These improvements contribute to a more sustainable and resilient core network, reducing energy consumption and lowering the overall carbon footprint, in alignment with global sustainability goals.

AI/ML capabilities are utilized across various domains in 5G, including the Network Data Analysis Function (NWDAF). The NWDAF is designed as a stand-alone functional entity within the 5G Core network architecture. It is intended to streamline the way core network data is produced and consumed, as well as to generate insights and take actions to enhance end-user experience.

NWDAF can be decomposed into three components: Analytics Logical Function (AnLF), the Model Training Logical Function (MTLF), and the Analytics Service [23].

2.2.1 NWDAF Energy Saving Actions

NWDAF may support different modes of operation and levels of reduced capability to enable energy savings. For example, NWDAF may advertise the various energy saving modes it supports by registering and updating its Network Function (NF) profile with relevant details. This profile may include information such as supported or unsupported analytics and statistics together with respective accuracy, along with their respective accuracy level. This allows consumers to subscribe to analytics they are interested in.

One possible action by NWDAF is advertising an energy cost per provided analytics. The energy cost per analytics may depend on the amount of data needed to train a model, storage required, analytics generation, and the actual cost of model training. The various factors affecting energy cost can be scaled with different weights to emphasize specific areas – e.g., the energy cost associated with handling large volumes of training data versus the energy hungry consumed in the actual training of the ML Model. Assigning weights to these factors helps prioritize energy optimization, balancing data processing with training costs.

The advertised energy costs may also be a complex parameter incorporating various cost factors instead of, or in addition to, the total energy costs. As part of the estimated costs, NWDAF can provide alternatives for modifying input parameters of the analytics request/subscription to make it more energy efficient without significantly reducing the accuracy.

The final decision to accept an alternative and modify the analytics request/subscription rests with the NF consumer based on internal logic, as NWDAF does not know how the consumer plans to use the requested analytics.

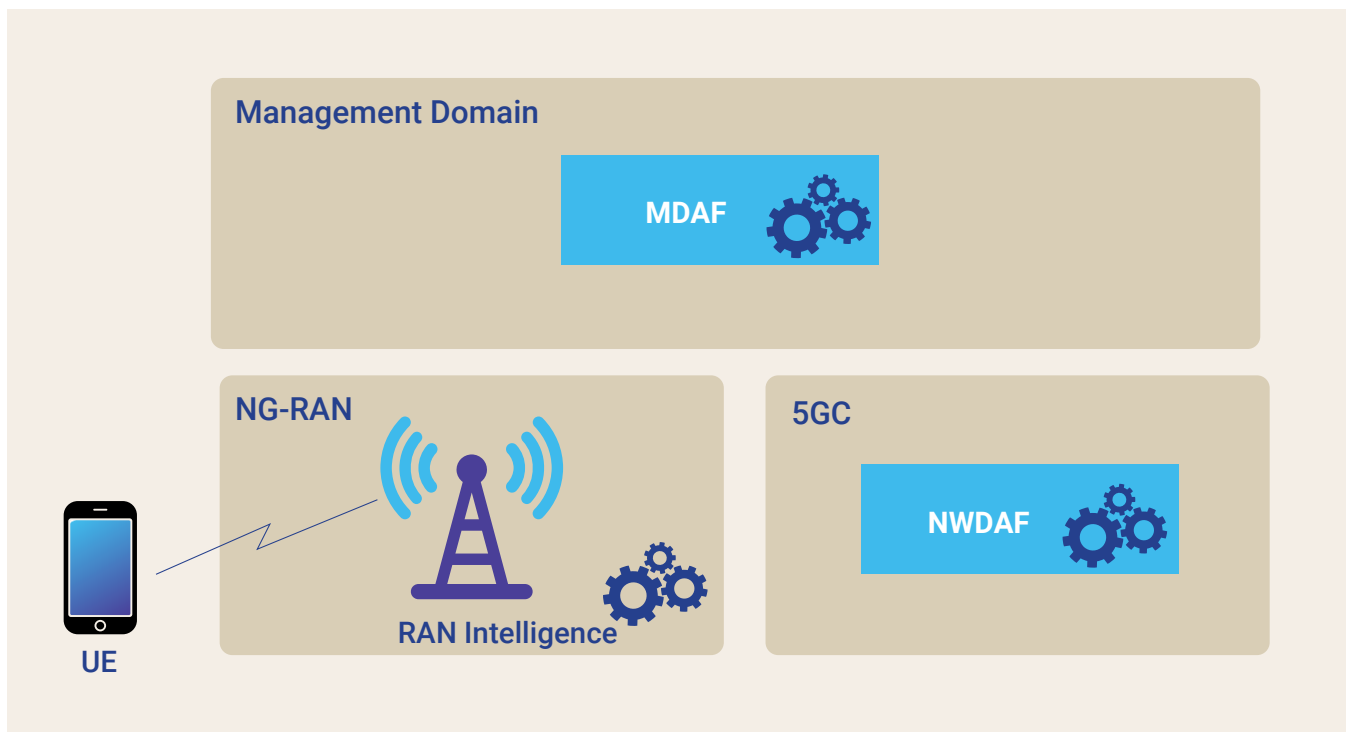


Figure 9: AI/ML in 5G System [24]

2.3 User Equipment

While a single mobile device consumes significantly lower energy than the RAN, the vast number of devices (including smartphones, tablets, and Internet of Things (IoT) devices) in mobile networks collectively consume a substantial amount of energy. Therefore, leveraging AI in a plethora of wireless devices as illustrated in Figure 10, can optimize power consumption for both application usage and connectivity.

Furthermore, when on-device AI, such as inference, for applications like AI-powered personal assistants and entertainment shifts processing from centralized cloud servers to the mobile device itself [25], the localized approach will minimize energy-intensive network connectivity, only utilizing it when necessary. As a result, the overall power consumption by devices and networks is significantly lowered, latency is reduced, and user experience can be improved.

To optimize the efficiency of on-device AI, computing resources like Central Processing Units (CPUs), Graphical Processing Units (GPUs), Neural Processing Units (NPU), and Tensor Processing Units (TPUs) are specifically tailored to AI operations and tasks. The targeted computing resource utilization maximizes both performance and energy efficiency, further contributing to a more sustainable mobile experience. More discussion on this subject can be found in the ATIS study on Generative AI in telecommunications [26].

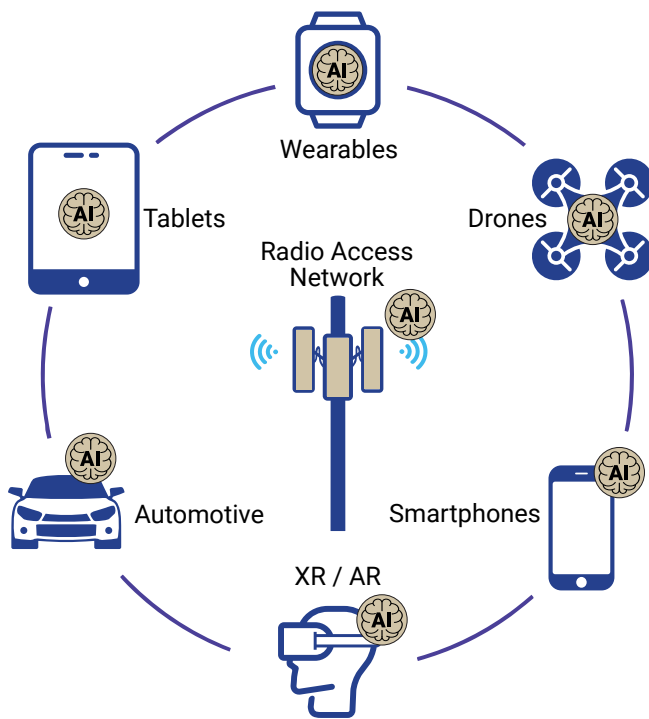


Figure 10: AI-Powered Mobile Devices for Power Consumption Optimization [27]

2.4 Data Center

Data centers process and store vast amounts of data, resulting in high energy consumption due to the need for constant power and cooling. Being the powerhouses of digital infrastructure, data centers stand to gain significantly from sustainable AI practices, which are essential for optimizing server utilization, improving cooling system management, and predicting maintenance needs. Implementing AI-driven energy management systems can lead to more precise control over the response to environmental condition variations within data centers, ensuring optimal resource utilization. This could result in significant reductions in energy consumption, environmental impact, and operational costs to achieve the sustainability targets.

The Next G Alliance (NGA) report “Evolution of Sustainability Indicators for Data Centers and Next Generation Core Networks,” [28] provides an overview of key data center sustainability strategies critical for achieving Net Zero emissions across the entire ICT value chain. Listed below is an excerpt from this report outlining sustainability attributes for data centers. There are multiple sustainability attributes that need to be considered when implementing a sustainability strategy to reduce the impact on the environment. The following is a list of multiple sustainability recommendations:

- > Energy efficiency: Data centers should optimize energy use to reduce carbon footprint and operational costs.
- > Renewable energy use: Incorporating renewable energy sources such as solar, wind, or hydro power to power data center operations.
- > Water conservation: Implementing water-efficient cooling systems and minimizing water usage for cooling purposes.
- > Waste reduction: Implementing strategies to reduce, reuse, and recycle waste generated by data center build and operations.
- > Carbon footprint: Minimizing GHG emissions associated with data center build and operations.
- > Green building design: Constructing data centers with environmentally friendly materials and energy-efficient designs.
- > Sustainable sourcing: Procuring equipment and materials from sustainable and ethical sources.
- > Environmental impact monitoring: Tracking and reporting environmental metrics to identify areas for improvement and measure progress over time. Mainly driven by the hyperscale’s, the data center industry has recognized that it must transform itself into incorporating sustainable operations.

2.4.1 Existing Data Center Sustainability Metrics

A holistic, yet simple framework to measure data center sustainability performance should include (at least) the following performance metrics:

Table 1: Data Center Sustainability Performance Metrics [28]

Energy Efficiency	Facility – Power Usage Effectiveness (PUE)
Zero Carbon	Renewable Energy Factor (REF)
Energy Reuse	Energy Reuse Factor (ERF)
Water Usage	Water Usage Effectiveness (WUE)

In summary, the NGA report [28] can serve as a guide for adopting a holistic approach to data center sustainability.

3 MEASURING THE ENVIRONMENTAL SUSTAINABILITY OF AI

Assessing the sustainability of AI in networks is crucial to ensuring that its benefits outweigh its environmental costs. This evaluation should cover factors such as energy consumption, environmental impact, and operational efficiency. Through AI-driven enhancements, such as predictive maintenance and performance optimization, operators can improve network efficiency and user experience. However, these benefits come with an increased energy demand due to the training, deployment, and updating of AI models, which necessitates careful management of resources.

This applies also to the environmental sustainability of AI systems. The environmental impact of the entire AI system life cycle is to be considered when evaluating its environmental sustainability. The advent of 6G necessitates a comprehensive paradigm shift. Figure 11 illustrates the three crucial dimensions in terms of implementing environmental sustainability metrics: Observability, Choice, and Circular Economy (OCC) [29].

Observability

Observability is vital for AI development, particularly in tracking and reporting environmental impacts throughout the AI model lifecycle, which demands significant resources. However, generating excessive data through observability can increase energy consumption, potentially outweighing the energy savings. Effective observability in AI helps balance performance and sustainability, optimizing resource use while navigating the trade-offs between accuracy and energy efficiency.

Choice

AI further empowers the concept of “choice” in telecom by enhancing energy efficiency and optimizing operations across various domains, including the RAN, Core Network, UE, and Data Centers. Implementing choice allows trade-off between the performance and sustainability metrics. Through advanced models, near real-time monitoring, and AI-driven recommendations, operators can fine-tune their network configurations, ensuring that performance enhancements are inextricably linked with sustainability. Gartner predicts that 30% of generative AI projects will be abandoned after proof of concept by the end of 2025 due to unclear business value [30], which makes the decision even more critical in determining when to use AI and when not to.

Circular Economy

Reaching Net Zero will require the telco industry to move from linear to circular business models. AI has the potential to accelerate the transition towards circular economy [31] by overcoming barriers to circularity, such as the low residual value of used products and high recycling costs. Some of the key strategies enabled by AI could include improving product utilization through extended lifecycles and product-as-a-service models.

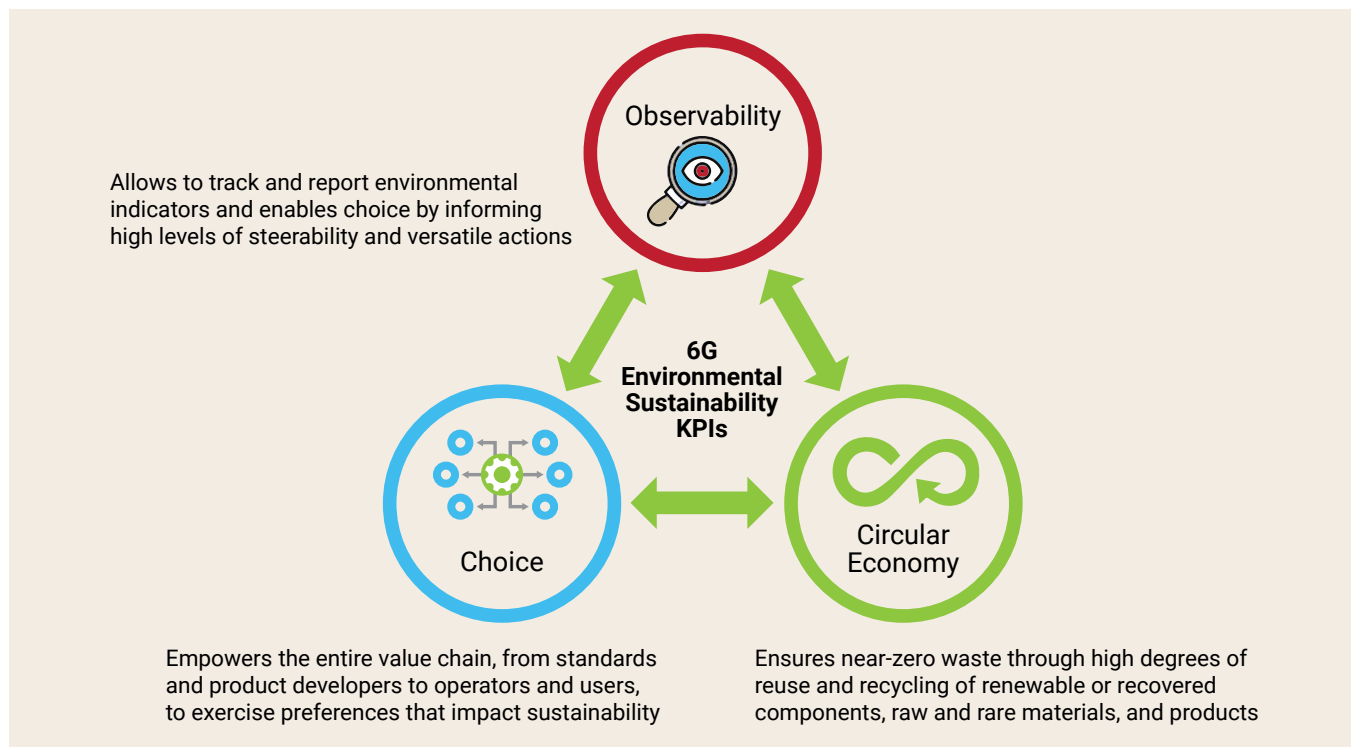


Figure 11: Enhanced Environmental Sustainability Framework

3.1 Measuring and Monitoring of AI/ML Energy Efficiency

To properly assess energy consumption during AI model training, network operators must employ a comprehensive strategy that includes tracking resources and evaluating hardware efficiency. This can be achieved by monitoring data center operations with energy meters and software (SW) tools like PUE, which measure energy efficiency. Techniques such as distributed training and energy-efficient algorithms can reduce environmental impact. Regular assessments will help identify areas for energy optimization [32].

Updating AI models as network conditions evolve, such as changes in traffic patterns or adjustments to topography, is also energy-intensive. Therefore, network operators should set a baseline for energy usage during these updates and weigh the energy consumed against the operational improvements achieved, such as reduced outages and better bandwidth management. This balance is essential to ensuring that AI strategies remain both environmentally sustainable and beneficial to network performance [33].

Energy consumption is one of the factors which significantly contribute to network OPEX and environmental impact. Thus, the energy consumption of the network is an important metric to monitor, control and reduce. The importance of optimizing network energy consumption has been ever increasing due to global sustainability requirements as well as increasing energy costs. In the future, accurate forecast and control of network energy consumption will be a key factor in adoption of new technologies.

AI/ML techniques have been employed in current networks, enabling the learning of systems states and providing recommendations on the actions to be taken. Such approach helps in efficient handling of a variety of network use cases (e.g., mobility management, traffic steering, resource optimization, etc.), AI/ML is envisioned to become a ubiquitous capability of the system, employed in all domains and planes of the network on much wider scale than employed today. Periodic retraining of the AI/ML models would be required to achieve best results coping up with the behavior of the system. This may imply significant energy consumption, in particular for large and complex models. Therefore, the impact of AI/ML on the energy consumption as well as how to handle the AI/ML in most energy efficient way are fundamental questions that need to be addressed. Furthermore, this may impose different requirements on the network system, especially in terms of capability of the system to enable monitoring of AI/ML energy consumption, as well as energy-aware management of AI/ML solutions – e.g., energy aware (re-)training, model discovery, and selection.

Measuring and monitoring the AI/ML energy consumption is a cornerstone for energy-aware management of AI/ML. There may be different approaches to express the measure of AI/ML energy consumption. The indication of the AI/ML energy consumption may be independent on the underlying platform that is running an AI/ML and may be expressed by means of complexity indicator of AI/ML. Such complexity

indication can be given in terms of Floating-Point Operations Per Second (FLOPs) – the number of floating-point operations needed for AI/ML execution. This metric quantifies the computational complexity of the algorithm itself – e.g., for executing inference without considering a certain architecture of a platform performing the operation. Such approach enables estimation of energy consumption of the models based on their complexity as well as comparison between different models. Another approach is expressing the AI/ML energy consumption in joules or watts, which is related to the specific execution platform and may enable a comparison of the energy consumption of AI/ML implementations on one specific platform.

The AI/ML energy consumption metrics may be monitored across different stages of the AI/ML life cycle (e.g., training and inference), enabling the application of energy-saving strategies tailored to each specific stage. Moreover, given the energy consumption requirements or constraints, the AI/ML life cycle stages and processes can be managed accordingly, e.g., training and re-training may be performed such that required maximum energy spending is respected.

The energy consumption of AI/ML solutions depends on a large number of factors, including those from different use cases, design, and the respective life cycle stages, such as model complexity, execution environment, required AI/ML performance, and amount of data needed for training/re-training/inference. For example, more complex solutions, trained on larger amount of data may show better performance but also consume far more energy compared to simpler approaches. Identifying and exposing such dependencies are prerequisites for performing optimization of AI/ML energy consumption. Such dependencies may be captured as part of the AI/ML energy consumption descriptor which indicates the AI/ML energy consumption subject to other related factors such as Key Performance Indicators (KPIs), execution time or speed, hardware and system specifications, etc. Such AI/ML energy consumption descriptor may provide sufficient information on the AI/ML energy profile to enable energy-aware registration, exposure, discovery and selection of AI/ML solutions meeting energy requirements and constraints.

3.1.1 Energy Efficiency of AI/ML Entities

AI/ML energy consumption measure gives the information on the energy needed to apply specific AI/ML enabled solutions and can be used for benchmarking of different solutions. However, energy consumption is often not a sufficient indicator to perform such benchmarking and trade-off analysis in fully comprehensive way. The fact that one solution consumes less energy than another does not necessarily imply that the solution is better even from energy point of view, as it can be highly inefficient – e.g., handling only small amount of load or being very inaccurate. Energy efficiency is a better KPI for comparing different AI/ML solutions addressing a given use case. Energy efficiency indicates the useful output per unit of energy consumed. For example, in the case of 5G core or RAN, it is often simplified as the ratio of data volume transferred to the energy

consumed. However, other useful outputs, such as coverage, the number of active subscribers served, and latency achieved, also contributed to the network's overall energy efficiency [34],[35].

This is why energy efficiency is never an absolute indicator, its meaning and value depend on the context and the system being reviewed:

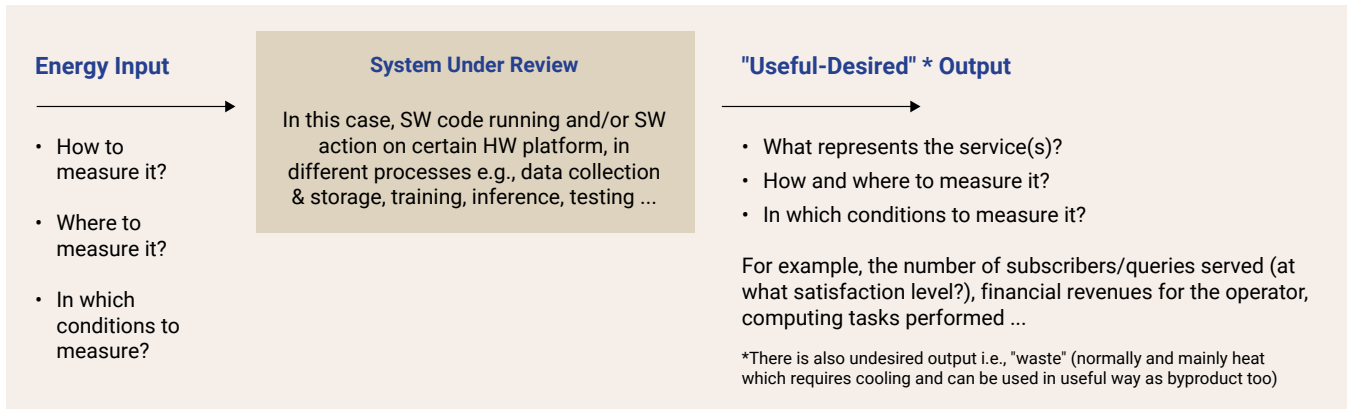


Figure 12: Energy Efficiency (EE): The Ratio of ["useful/desired" output] / [energy input] of a Given System

The energy efficiency of an AI/ML entity over the entire life cycle can be characterized via below examples:

Table 2: Energy Efficiency of AI/ML Processes Throughout the Life Cycle

The Different AI/ML Processes	Examples of Possible EE metrics	All denominators are the measurement of energy used (unit in Watt-hours or Joule) during the respective processes.
Data Collection & Pre-Processing	$EE_{AI\ DC\ \&\ P} = \frac{f(\text{data collection, storage, processing})}{\text{Total energy used during the process}}$	The numerator may represent the amount of relevant data collected and stored for the training and pre-trained according to the training needs, which could be in terms of data volume (e.g., gigabytes) or another relevant data volume unit.
Training	$EE_{AI\ TR} = \frac{f(\text{performance indicators after training})}{\text{Total energy used during the process}}$	The numerator may represent the performance achieved by the trained model, such as accuracy, precision, recall, F1 score, or any other relevant metric.
Testing	$EE_{AI\ TE} = \frac{f(\text{testing performance indicators})}{\text{Total energy used during the process}}$	The numerator may represent the AI/ML testing performance/ effectiveness on the test data, such as accuracy, precision, recall, F1 score, or any other relevant metric.
Inference	$EE_{AI\ I} = \frac{f(\text{inference performance indicators})}{\text{Total energy used during the process}}$	The numerator may represent the performance achieved during the inference process, such as throughput (the number of inferences per second), latency (execution time of an inference), accuracy, precision, recall, F1 score, latency, or any other relevant metric.
All Processes Included	$EE_{AI} = \frac{f(\text{performance \& quality indicators})}{\text{Total energy used during the process}}$	The numerator may represent the function of the relevant performance achieved by the AI/ML entity, such as its throughput, its desired output. Depending on the use case and its context, it could consider one or a set of performance and quality indicators (e.g. the number of inferences per second), latency (execution time of an inference), accuracy, precision, recall, F1 score, latency, or any other relevant metric.

4 TECHNIQUES FOR REDUCING ENERGY CONSUMPTION IN MACHINE LEARNING

AI offers promising solutions to enhance the energy efficiency of these networks by optimizing various aspects such as network resource allocation, load balancing, and predictive maintenance. However, the implementation of AI itself can be energy intensive. To address this, several AI techniques can be employed to minimize energy consumption, including the use of energy efficient algorithms, model compression, edge computing, and federated learning. These techniques not only reduce the environmental footprint of AI, but also ensure that the deployment of AI in mobile networks contributes positively to overall energy efficiency goals.

Reducing CPU and energy consumption during AI training and inference is crucial for making AI more sustainable. Here are some key techniques that can be applied:

4.1 Model Compression Techniques

4.1.1 Pruning

This technique involves removing redundant or less important parameters from a model, reducing its size and computational requirements. Pruned models are lighter and faster, leading to reduced energy consumption during both training and inference. Pruning involves removing unnecessary weight connections to address over-parameterization. Inspired by the human brain's pruning process, this technique includes unstructured pruning, which removes individual weight connections and relies on hardware's ability to handle sparse operations, and structured pruning, which removes groups of connections such as channels or filters, resulting in higher energy efficiency. The pruning process starts with a dense pre-trained network, where weights are evaluated and removed based on specific criteria, followed by fine-tuning the remaining weights. Methods include saliency (Hessian-aware), which measures the sensitivity of training loss to each weight; magnitude-based, which removes the smallest weights first; and energy-aware, which targets weights consuming the most energy to maximize savings.

4.1.2 Quantization

Quantization reduces the precision of the model parameters (e.g., from 32-bit floating point to 8-bit integers), which can significantly decrease the computational load and power usage without substantially sacrificing accuracy. It reduces the bit precision of model weights and activations, lowering both memory requirements and computational load. There are different types of quantization: symmetric quantization, often used for weights with a fixed clipping range; dynamic quantization, used for activations with a dynamically computed clipping range; and static quantization, which uses a calibrated, fixed clipping range from a representative dataset. Mixed-precision quantization can be applied to

enhance performance further. The primary methods for quantization are Post-Training Quantization (PTQ) and Quantization-Aware Training (QAT). PTQ quantizes weights and activations after training without fine-tuning, making it a quick but potentially less accurate method. QAT incorporates quantization during training to maintain accuracy, using floating-point precision during gradient updates.

4.1.3 Knowledge Distillation (KD)

In this approach, a smaller model (student) is trained to mimic the outputs of a larger, pre-trained model (teacher). The smaller model requires less computational power and is more energy-efficient during inference. KD compresses a large model into a smaller one by transferring knowledge from one or more computationally expensive teacher models to a single, smaller student model. This process aims to make the student model replicate the teacher model's behavior, thus improving accuracy. Importantly, the student model can have a different architecture from the teacher model, allowing for greater flexibility and potential performance gains.

In KD, the class probability of an input, given the logits from a network, is calculated using the softmax function with a temperature parameter, T :

$$p_i = \frac{\exp(\frac{z_i}{T})}{\sum_j \exp(\frac{z_j}{T})}$$

Here T , controls the smoothness of the probability distribution. When $T = 1$, this is the standard softmax function. For $T > 1$, the probabilities become more uniform, a process known as softening the probability distribution. Softened distributions are crucial in KD as they provide richer information than hard class labels, aiding the student model in learning from the teacher model's outputs.

The training loss in KD is: $L = a \times L_{student} + L_{KD}$

where $L_{student}$ is the student model's loss, and a is a parameter that balances the losses. L_{KD} represents the *Kullback-Leibler* (KL) divergence between the student (p) and teacher (q) models' probability distributions.

The core idea of KD is that soft labels from a well-trained teacher network contain more information about a data point than hard labels. High probabilities for multiple classes indicate that the input's class is near the decision boundary. For example, in digit classification, an input "1" would have a higher soft probability for class "7" than for "8" since "7" resembles "1" more. Training a student network to match these probabilities allows it to gain insights from the teacher beyond what hard labels convey.

Several techniques extend the original concept of KD. The *Teacher Assistant* method suggests that for a given student size, there is an optimal teacher size for maximum accuracy.

If the teacher model is too similar in size to the student, knowledge transfer may be ineffective. Conversely, if the size gap is too large, the student may struggle to mimic the teacher's complex behavior. An optimal teacher can be refined by a larger teacher, with intermediate teachers acting as assistants.

Another technique involves training deeper but thinner student networks, which can effectively learn complex representations with a fixed number of operations. **Attention transfer** uses attention maps from multiple intermediate teacher layers as hints for the student, highlighting input areas most influential in output decisions.

In some cases, using not only the teacher's logits but also the function's derivatives to train the student can capture more information about the target function. This method, known as **Sobolev Training**, can improve the student's performance. Finally, using multiple teachers to train a single student, called **Oracle Knowledge Distillation**, can enhance performance by averaging the logits of correct predictions from several teachers, similar to ensemble methods.

To illustrate the effectiveness of KD, Figure 14 presents a comparative curve showing the results of applying this technique on DeepRx [36] versus training from scratch.

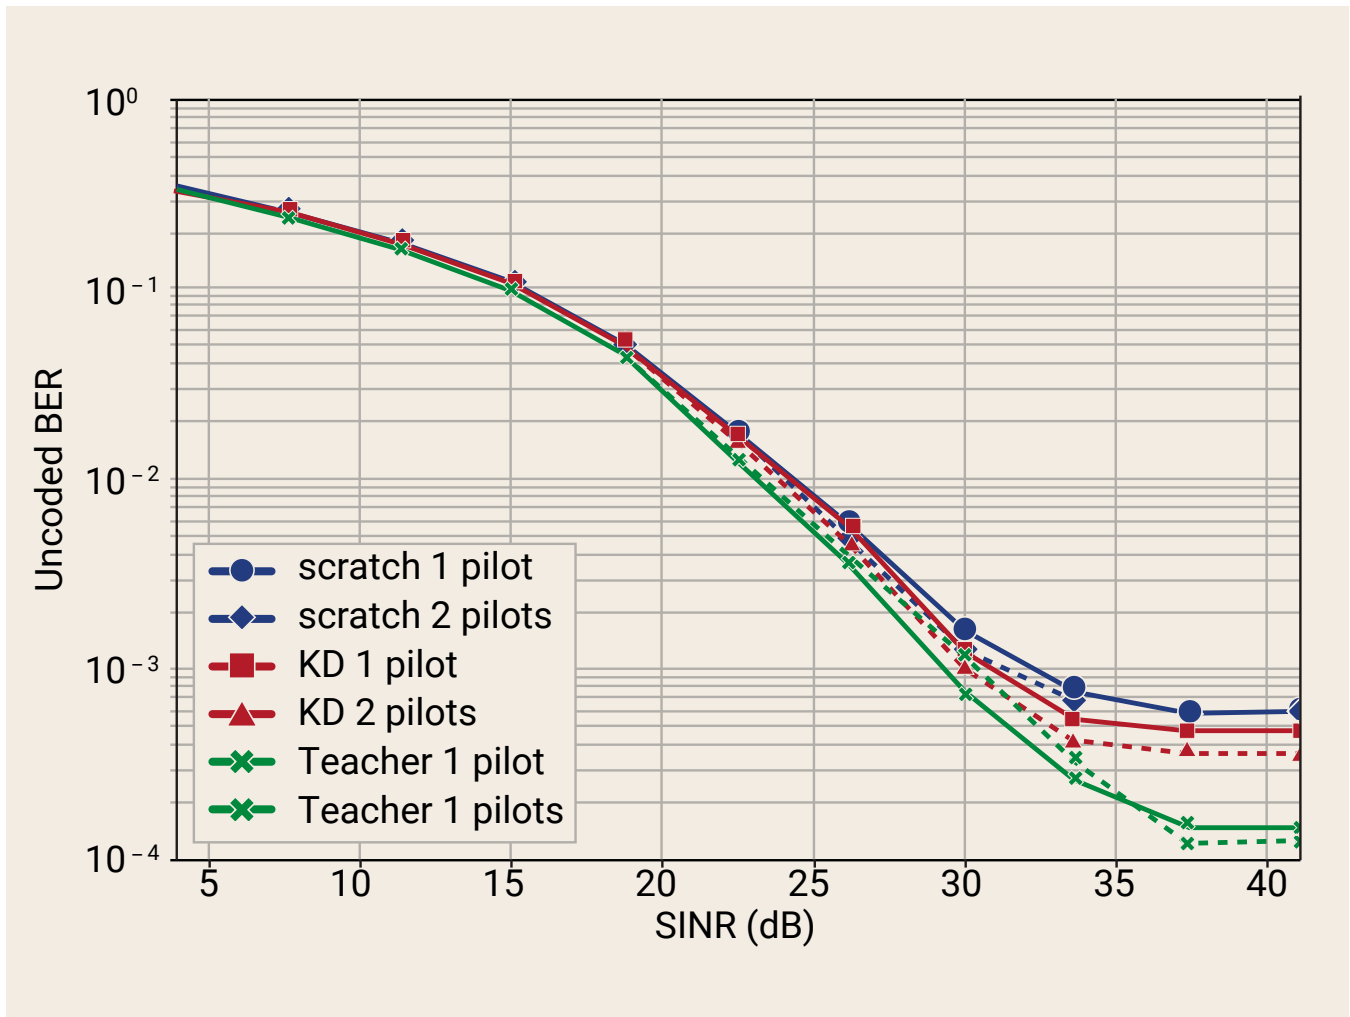


Figure 14: Comparison of Bit Error Rate (BER) Performance for DeepRx Models Using Knowledge Distillation Versus Training from Scratch [36]

4.2

Additional Techniques

4.2.1 Efficient Architectures

Low-Complexity Neural Networks: Using architectures specifically designed for efficiency, such as MobileNets, EfficientNet, or SqueezeNet, can greatly reduce the energy required for both training and inference.

Neural Architecture Search (NAS): NAS can be employed to automatically design neural network architectures that optimize for both performance and efficiency, leading to models that are less resource-intensive. It automates the discovery of the most effective neural network architecture. This involves exploring various architectural choices to balance accuracy and energy efficiency. Optimization algorithms are used to fine-tune the model iteratively. To further reduce energy consumption, it is advisable to limit epochs and data volume and apply early stopping when no performance improvement is observed.

4.2.2 Gradient Accumulation

During training, gradient accumulation allows the use of smaller batches while achieving the same effect as larger batches. This reduces memory usage and can lower the energy requirements of training, especially on resource-constrained devices.

4.2.3 Sparse Representations

Sparse representations involve structuring the model in a way that most of the parameters are zero, which reduces the number of operations required for training and inference. Techniques like sparse regularization during training can promote sparsity in the model.

4.2.4 Adaptive Computation

Dynamic Inference: Adjusting the computation during inference based on the input data can save energy. For example, models can be designed to stop computation early if a confident prediction is made with less computation, reducing unnecessary operations.

Conditional Computation: This technique allows parts of the network to be activated only when needed, thus reducing the overall computational load.

4.2.5 Edge Computing

Performing AI inference closer to the data source (e.g., on edge devices) can reduce the need for transmitting data to centralized servers, thereby saving energy. Edge AI models are typically lightweight and optimized for lower energy consumption.

4.2.6 Federated Learning

Federated learning distributes the training process across multiple devices, keeping the data local and reducing the need for data transfer, which can lower energy consumption associated with data movement and centralized processing.

4.2.7 Efficient Hardware Utilization

Utilizing Specialized Hardware: Leveraging hardware accelerators such as GPUs, TPUs, or AI-specific chips like Google's Edge TPU can lead to more energy-efficient processing compared to general-purpose CPUs.

Parallelization: Efficiently parallelizing tasks and fully utilizing hardware resources can reduce the time and energy required for training and inference.

4.2.8 Energy-Aware Scheduling

Implementing scheduling algorithms that are aware of the energy consumption of different tasks can optimize the deployment of AI models, ensuring that energy-intensive operations are minimized during peak energy usage times.

4.2.9 Emerging AI Technology

Low-power, nature-inspired AI/ML architectures offer both high accuracy and compact design. These models consume less energy due to their sparsity, as fewer connections are required than the maximum possible number of connections within that network between each layer, thereby reducing computational demands. Additionally, the model size is minimized since there are significantly fewer parameters to train [37].

By integrating these techniques, AI systems can achieve a balance between high performance and low energy consumption, contributing to more sustainable and environmentally friendly AI applications.

4.3

Energy Optimization Techniques for Reducing AI/ML Training and Inference Costs

The training and inference phases are among the most energy-intensive stages of the AI/ML life cycle. Energy consumption during training scales with the size of the dataset, the complexity of the model architecture, and the number of parameters involved. The amount of computational power used to train the largest AI models had doubled every 3.4 months since 2012 [38]. As models grow larger, the energy demands of these processes increase exponentially, making energy optimization critical.

The energy consumption of the inference process, while negligible compared to training for a single inference, can vary significantly depending on the use case. In scenarios with intensive and frequent inference, the cumulative energy usage of inference can surpass that of training, making it a critical factor in overall energy optimization.

Energy optimization techniques applied to AI/ML training and inference are crucial, as they can significantly reduce energy consumption. Even though these optimization techniques may impact AI/ML performance, such as causing slight decreases in accuracy or longer training times, these trade-offs are often acceptable, especially given the benefits in terms of energy efficiency. For example, hyperparameter

tuning techniques can help strike the right balance between performance and energy consumption by selecting optimal training parameters. This can lead to substantial energy savings, even if it requires accepting a slight drop in model accuracy.

4.4 The Importance of Key Performance Indicators (KPIs) and Key Value Indicators (KVIs)

As AI becomes increasingly integrated into cellular networks, understanding its environmental impact is crucial. KPIs and KVIs are critical for assessing the sustainability impact of AI technologies. KPIs are quantifiable measures used to evaluate the efficiency and performance of specific activities, such as the energy consumption of data centers or the carbon footprint of training models. They provide a clear metric to gauge progress toward sustainability goals. For instance, energy efficiency can be measured in kilowatt-hours per task, while carbon emissions can be tracked in CO₂ equivalents. On the other hand, KVIs focus on the value-driven aspects, reflecting broader, qualitative outcomes like social and environmental impacts. These may include metrics like the reduction in resource depletion, improvements in biodiversity conservation through AI applications, or the enhancement of human well-being by reducing hazardous emissions.

Defining KPIs and KVIs requires a comprehensive understanding of the AI lifecycle and the associated environmental and social impacts. It's essential to consider both direct factors, such as electricity usage and cooling needs of data centers, and indirect factors, like the sourcing of rare earth materials for hardware. Incorporating KVIs into sustainability assessments ensures a holistic approach, emphasizing not only efficiency but also the broader values and ethical considerations of AI deployments [33],[39].

In the telecom sector, KPIs and KVIs play an important role in monitoring and enhancing sustainability practices. KPIs in this domain often focus on the energy efficiency of network infrastructure, such as base stations and data centers, as well as the carbon footprint associated with the deployment and operation of telecommunication networks. Metrics like energy consumption per gigabyte of data transmitted and the proportion of renewable energy used in operations are critical for assessing sustainability. Meanwhile, KVIs in telecom encompass broader, value-driven measures such as the socio-economic benefits of expanding network coverage to underserved areas, which can improve access to education and healthcare. Additionally, the environmental impact of electronic waste from outdated equipment and the adoption of circular economy practices, such as recycling and reusing components, are key considerations. These indicators help telecom companies balance operational efficiency with environmental stewardship and social responsibility, guiding efforts towards a more sustainable and inclusive digital future [40].

AI algorithms, especially deep learning models, are notorious for their high energy consumption. Training a large-scale

model, such as GPT-3, can require as much electricity as the yearly consumption of several households [32].

Factors such as model complexity, hardware efficiency, and data center energy sources significantly impact these figures. Efforts to improve energy efficiency include optimizing algorithms, using specialized hardware like TPUs, and shifting towards renewable energy-powered data centers [41]. The sustainability impact of AI extends beyond energy consumption during training and inference. The entire lifecycle, from data collection and preprocessing to hardware manufacturing and disposal, plays a role.

4.5 Trade-off Between Energy Consumed Versus Energy Saved

To reduce the computational load of DeepRx, the size of the central layers was decreased. By employing KD, a reduction in energy consumption can be achieved. The student model, with a size of 11 Trillion Floating-Point Operations Per Second (TFLOPs), consumes significantly less energy compared to a 30 TFLOPs model. This reduction in size translates to lower computational and memory requirements, making it more efficient for deployment in mobile networks. The 11 TFLOPs student model uses approximately one-third of the computational resources compared to the 30 TFLOPs teacher model, resulting in substantial energy savings. This efficiency is particularly beneficial in mobile environments where power and memory resources are limited.

Table 3: Energy Consumed vs Energy Saved

Model Type	Size (TFLOPs)	Energy Consumption (Relative)	Energy Savings (Relative)
Teacher Model	30	High (3x)	Baseline
Student Model	11	Low (1x)	Significant (2x reduction)

Teacher Model: With a size of 30 TFLOPs, the energy consumption is relatively high, serving as the baseline for comparison.

Student Model: With a size of 11 TFLOPs, the energy consumption is significantly lower, using approximately one-third of the computational resources, leading to substantial energy savings.

By comparing the teacher and student models, a clear trade-off is observed between the size of the model and the energy consumed, with the student model offering significant energy savings at the cost of (controlled) reduced performance. This makes the student model particularly advantageous for deployment in mobile networks where energy efficiency is crucial.

5 FUTURE LOOKING ASPECTS

The emerging era of telecommunications demands a high focus on both performance and sustainability. AI-driven techniques are emerging as essential tools to optimize network operations while minimizing environmental impacts. From intent-driven automation to AI/ML integration at various network layers, AI is set to revolutionize how MNOs balance service delivery with energy efficiency. The introduction of quantum computing further accelerates the potential for sustainable networks by drastically enhancing computational efficiency. Non-Terrestrial Networks (NTNs) are also a key part of the 6G landscape, as they extend coverage to remote areas, and integrating AI into NTNs offers new ways to enhance the sustainability and operational efficiency of satellite communications. Together, these innovations form a comprehensive strategy for building networks that not only deliver high-quality services but also support global efforts toward environmental responsibility.

5.1 AI/ML and the Physical Layer

AI is anticipated to enhance the spectral efficiency of the 6G physical layer, potentially improving the overall energy efficiency of radio networks. Early proofs-of-concept demonstrate the performance benefits of ML-based transmission and reception at the physical layer. However, further investigation is needed to determine the hardware requirements for efficient neural network inference and to standardize aspects of 6G that will fully utilize AI at this level. In particular, achieving energy-efficient real-time inference of physical-layer ML models may necessitate specialized neural hardware accelerators, as conventional GPUs and general-purpose processors are not energy-efficient enough.

Moreover, the upper layers of the radio protocol stack primarily rely on heuristic algorithms, which differ from the solid information-theory foundations of physical-layer solutions. The absence of provable solutions to complex problems, such as scheduling, offers opportunities to discover innovative solutions through data-driven approaches. The less stringent runtime requirements of these upper layers make them ideal candidates for neural network deployment. For 6G to effectively support ML-based protocol solutions, it must be designed from the ground up, incorporating robust Machine Learning Operations (MLOps) procedures and the flexibility to deploy both traditional expert systems and machine-learned logic. This approach will enable the development of highly tailored, energy-efficient, and context-aware protocol stacks, with promising research areas in 2024 including ML-based Medium Access Control (MAC) protocols, Discontinuous Reception (DRX) control, Downlink Control Information (DCI) compression, resource scheduling, and power control.

5.2 AI Intent-Driven Techniques in 6G

In the future, AI-driven sustainability techniques will play a central role in achieving the environmental and business goals of network operators. These techniques will enable intent-driven automation, where AI continuously optimizes network resources to meet specific objectives, such as reducing energy consumption or lowering carbon emissions. Similar to self-driving cars, various automation functions such as adaptive cruise control and lane-keeping, telecom networks will evolve towards full autonomy. AI systems will manage everything from traffic routing to energy-efficient resource allocation, reducing the need for manual intervention and optimizing performance with sustainability as a key factor.

In this context, “self-driving” or intent-based networks will interpret high-level objectives from network operators, such as delivering high-quality video streaming to Fixed Wireless Access (FWA) subscribers while minimizing power consumption. The network will autonomously adjust key parameters, such as bandwidth allocation, data traffic prioritization, and energy-saving mechanisms, to meet these goals. This approach ensures that energy-efficient practices are seamlessly integrated into every aspect of network operations. As 6G and future systems develop, AI-enabled sustainability will be embedded into the network’s architecture, enabling networks to dynamically adapt and deliver services with the least environmental impact, ensuring alignment with sustainability requirements at every level.

5.3 Leveraging Quantum AI for More Sustainable and Resilient Mobile Networks

Quantum computers hold the potential for superior efficiency compared to classical computers for specific computations due to their unique characteristics. This efficiency arises from quantum computers’ ability to utilize quantum mechanics principles, such as superposition, entanglement, and quantum parallelism, allowing them to execute complex calculations significantly faster than traditional systems. This increased efficiency can be particularly beneficial for sustainable AI applications in several ways.

For instance, quantum computers can provide exponential speedup for solving complex optimization and data analysis problems, dramatically decreasing the time and energy needed for AI tasks. Additionally, their capability to handle high-dimensional data efficiently facilitates faster and more energy-efficient pattern recognition and predictions, which are essential for sustainable AI applications. Quantum algorithms can optimize resource allocation more effectively than classical methods, improving energy efficiency in applications such as energy grid optimization and logistics.

The benefits can also be seen in reduction of training times for ML models, making AI development more sustainable and resource efficient. Recent trials in Japan [42] have successfully utilized quantum computing to optimize mobile network paging processes, enhancing spectrum efficiency, and lowering energy consumption, showcasing a potential 15% improvement in reducing congestion during peak calling times.

5.4 AI Sustainability for Non-Terrestrial Networks

AI is revolutionizing the management and operation of NTN, particularly in satellite systems. By autonomously handling tasks such as orbit adjustments and power management, AI reduces the need for human intervention, allowing satellites to make real-time decisions that enhance operational efficiency and reliability. Additionally, AI plays a key role in predictive maintenance, where ML algorithms analyze satellite data to identify potential performance issues, enabling proactive maintenance and extending the lifespan of satellites. Beyond maintenance, AI will be integral to the coordination of satellites in 6G, optimizing coverage, minimizing interference, and ensuring seamless communication across vast satellite constellations [43]. As NTN become crucial in next-generation telecom systems, AI will be essential in managing these networks for greater efficiency and sustainability.

Organizations like the ITU are leading efforts to enhance the sustainability of NTN, with a focus on reducing carbon footprint. The environmental implications of Low Earth Orbit (LEO) satellites are currently under scrutiny, particularly concerning the increasing frequency of launches [44]. Strategies such as launching multiple satellites simultaneously are being explored to mitigate carbon emissions. Furthermore, the energy consumption of AI processing in space versus ground-based systems presents a significant contrast, as onboard processing is limited by power availability from solar panels, whereas ground-based systems often demand substantial energy for heavy AI tasks. This raises important considerations for energy efficiency as NTN continue to scale with larger constellations and more advanced AI applications. In general, the total energy consumption of an AI system in NTN is difficult to quantify precisely without specific data. However, given the energy-intensive nature of AI, the footprint could be significant, especially as NTN scale up with larger constellations and more complex AI applications.

SUMMARY, CHALLENGES, & RECOMMENDATIONS

The intersection of sustainable AI and AI for sustainability highlights the intricate balance required to achieve energy efficiency in telecommunications networks. AI models are increasingly employed to optimize energy consumption by managing hardware components, determining sleep cycle durations, and controlling server operations. However, the development and training of these AI models necessitate significant computational power and energy, which must be accounted for when assessing the net energy efficiency gains. Establishing a comprehensive framework that systematically evaluates energy consumption—the frequency of (re)training, the duration of application of energy efficiency actions, the number of sites over which the energy efficiency actions need to be applied, the cooling availability for the CPUs and GPUs to training the models, and the costs of energy.

6.1 Challenges

6.1.1 Emissions

The energy consumption of data centers contributes significantly to global electricity demand, accounting for approximately 1%. Without intervention, this figure is expected to rise exponentially, resulting in carbon emissions from data centers that rely on non-renewable energy sources. AI/ML algorithms can help optimize power usage by dynamically balancing workloads across servers and adjusting power and cooling requirements, thereby reducing overall energy consumption. Predictive maintenance enabled by AI can also identify hardware failures before they occur, minimizing energy consumption spikes.

6.1.2 Water Usage

Traditional data center cooling systems often rely heavily on water, which can strain local resources, particularly in water-scarce regions. AI can enhance water efficiency by fine-tuning cooling systems in real-time, adjusting water flow and temperature based on sensor data. For instance, Microsoft has implemented AI-driven cooling techniques that reduce water usage by up to 90% in specific data centers. AI models can also predict cooling needs based on weather forecasts, ensuring that water is used efficiently. Additionally, while air cooling systems are more eco-friendly, they require careful management to maintain energy efficiency, and AI can optimize airflow and fan speeds accordingly.

6.1.3 Waste Management

Data centers generate significant amounts of electronic waste from outdated hardware, which can contain harmful substances like lead, mercury, and cadmium. Effective waste management and hardware lifecycle oversight are critical to mitigate the environmental impact of this electronic waste.

By implementing AI-driven strategies, telecom networks can better manage their equipment lifecycles, ensuring responsible disposal and recycling of electronic components to minimize harmful waste.

6.2 Recommendations

To enhance AI-driven sustainability in telecom networks, several key recommendations can be adopted. AI algorithms can analyze user behavior, network conditions, and environmental data to propose service configurations that align with users' sustainability preferences while ensuring satisfactory Quality of Service (QoS) [29]. KPIs such as the effectiveness of sustainability-focused service recommendations and the improvement in energy efficiency through AI optimization should be defined and measured. Real-world applications can help track acceptance rates for AI-generated suggestions and quantify energy savings attributed to AI-powered enhancements. Addressing sustainability challenges, such as emissions reduction, water usage optimization, and waste management in data centers, through targeted AI and ML solutions will be crucial for promoting environmental stewardship in the next generation of wireless technology.

Reducing embodied carbon emissions in AI systems requires focusing on the full lifecycle of computing hardware. Key strategies include:

Sustainable Manufacturing: Use low-carbon materials and energy-efficient processes in hardware production to reduce the initial carbon footprint.

Extending Equipment Lifespan: Maximize hardware usage through maintenance and upgrades, minimizing the need for frequent replacements and associated emissions.

Circular Economy Practices: Promote refurbishing, reusing, and recycling of hardware to lower the demand for new materials.

Supply Chain Optimization: Work with suppliers to reduce emissions across the value chain, including sourcing from low-carbon regions and optimizing transportation.

Energy-Efficient Hardware Design: Focus on designing energy-efficient hardware, reducing both embodied emissions during manufacturing and operational emissions.

Adopting these strategies can significantly reduce the embodied carbon footprint in AI systems, supporting more sustainable practices.

Regarding standardization, ISO is in the process of finalizing technical report (TR) 20226, "Environmental sustainability aspects of AI systems" [45]. Similarly, European Committee

for Standardization (CEN)/ European Committee for Electrotechnical Standardization (CENELEC) has been developing a TR focused on environmentally sustainable AI. In addition, ITU-T has recently initiated efforts to standardize “Guidelines for Assessing the Impact of Artificial Intelligence on Greenhouse gas emissions” [46]. A standardized approach to this issue is a need of the hour. In terms of energy efficiency, ITU-T is also working on standards such as the “Deep Learning Computation Energy Efficiency Evaluation Framework and Metrics” [47] and the “Recommendation for the Design of Environmentally Sustainable AI-Based and XR-Based Systems” [48].

The U.S. White House Office of Science and Technology Policy (OSTP) has published the “Blueprint for an AI Bill of Rights,” [49] which outlines principles and best practices for the design, use, and deployment of AI systems to protect American rights in the AI era. Although it includes crucial aspects like protection against algorithmic discrimination and data privacy, it does not address the environmental sustainability of AI systems. However, the “Artificial Intelligence Environmental Impacts Act of 2024,” [50] introduced in the U.S. Senate on February 1, 2024, mandates a comprehensive study on AI’s environmental impacts, focusing on various stages of its lifecycle. The study will analyze energy consumption, pollution from AI models, the environmental footprint of AI hardware (including raw material extraction, manufacturing, and e-waste), and energy and water usage for cooling data centers. It will also assess how design and operational choices affect environmental outcomes and identify local impacts such as grid stress and water usage.

Additionally, the study will explore both the positive environmental benefits of AI, such as optimizing energy efficiency and aiding renewable energy development, as well as potential negative impacts like rebound effects and pollution from high-energy activities. Regional disparities in the environmental footprint of AI will also be considered. These initiatives from the U.S. government underscore the growing need to understand and manage the social and environmental impacts of AI.

AI holds immense promises for improving the environmental sustainability of next-generation networks by optimizing energy use, enhancing resource efficiency, and integrating renewable power sources. However, caution must be exercised in the implementation of AI to ensure that its own energy demands do not compound the problem. As AI becomes more embedded in network operations, it must be carefully managed to align with our Net Zero objectives, ensuring that its profound potential supports, rather than hinders, our transition to a more sustainable future.

7

ABBREVIATIONS AND ACRONYMS

AI	Artificial Intelligence
AnLF	Analytics Logical Function
BLOOM	BigScience Large Open-science Open-access Multilingual Language Model
CEN	European Committee for Standardization
CENELEC	European Committee for Electrotechnical Standardization
CPU	Central Processing Unit
DCI	Downlink Control Information
DOE	Department of Energy
DRX	Discontinuous Reception
EE	Energy Efficiency
ERF	Energy Reuse Factor
FLOPs	Floating-point Operations per Second
FWA	Fixed Wireless Access
GHG	Greenhouse Gases
GPU	Graphical Processing Unit
ICT	Information and Communications Technology
IoT	Internet of Things
KD	Knowledge Distillation
KL	Kullback-Leibler
KPI	Key Performance Indicators
KVI	Key Value Indicators
LCA	Life-cycle Assessment
LEO	Low Earth Orbit
LLMs	Large Language Models
MAC	Medium Access Control
ML	Machine Learning
MLOps	Machine Learning Operations
MNOs	Mobile Network Operators
MTLF	Model Training Logical Function
NAS	Neural Architecture Search
NF	Network Function
NPU	Neural Processing Unit
NTNs	Non-Terrestrial Networks
NWDAF	Network Data Analysis Function

8 REFERENCES

1. B. Washburn, "Omdia's new AI network traffic forecast expects surge, pressure on telcos to meet future growth," Omdia, Dec. 2023. <https://omdia.tech.informa.com/om119985/omdias-new-ai-network-traffic-forecast-expects-surge-pressure-on-telcos-to-meet-future-growth>
2. K. Hao, "Training a single AI model can emit as much carbon as five cars in their lifetimes," MIT Technology Review. <https://www.technologyreview.com/2019/06/06/239031/training-a-single-ai-model-can-emit-as-much-carbon-as-five-cars-in-their-lifetimes/>
3. M. Farzan and S. Kallio, "A transparent and standards-based way to assess the environmental impact of AI systems," Nokia, <https://onestore.nokia.com/asset/214115>
4. ISO/IEC 5338:2023, "Information technology – Artificial intelligence – AI system life cycle processes," 2023.
5. ISO 14040:2006, "Environmental management – Life cycle assessment – Principles and Framework," 2022.
6. ITU-T L.1410, "Methodology for environmental life cycle assessments of information and communication technology goods, networks and services," <https://handle.itu.int/11.1002/1000/16010>
7. World Economic Forum, "How to manage AI's energy demand – today, tomorrow and in the future." <https://www.weforum.org/agenda/2024/04/how-to-manage-ais-energy-demand-today-tomorrow-and-in-the-future/>
8. D. Kerr, "AI brings soaring emissions for Google and Microsoft, a major contributor to climate change," NPR, 2024. <https://www.npr.org/2024/07/12/g-s1-9545/ai-brings-soaring-emissions-for-google-and-microsoft-a-major-contributor-to-climate-change>
9. U.S. Energy Information Administration, "Electricity in the U.S.," Mar. 2024. <https://www.eia.gov/energyexplained/electricity/electricity-in-the-us.php>
10. RMI, "Clean Repowering Can Avoid Disastrous Delays with 250 GW of Clean Energy," June 2024. <https://rmi.org/clean-repowering-can-avoid-disastrous-delays-with-250-gw-of-clean-energy/>
11. Western Public Radio, "Powering Intelligence: Analyzing Artificial Intelligence and Data Center Energy Consumption," May 2024. https://www.wpr.org/wp-content/uploads/2024/06/3002028905_Powering-Intelligence_-_Analyzing-Artificial-Intelligence-and-Data-Center-Energy-Consumption.pdf
12. EY and Liberty Global, "Wired for AI." https://assets.ey.com/content/dam/ey-sites/ey-com/en_uk/topics/ai/wired-for-ai-ey-and-liberty-global-report.pdf
13. MCE Clean Energy, "How peaker plants impact California communities," Sep. 2022. <https://mcecleanenergy.org/how-peaker-plants-impact-california-communities-2/>
14. California ISO, "Historical CO2 Trend," <https://www.caiso.com/todays-outlook/emissions#section-historical-co2-trend>
15. U.S. Environmental Protection Agency, "Data Explorer," <https://www.epa.gov/egrid/data-explorer>
16. GSMA Intelligence, "Going Green: benchmarking the energy efficiency of mobile networks," Feb. 2023. <https://data.gsmaintelligence.com/api-web/v2/research-file-download?id=74384072&file=280223-Going-Green-Second-Edition.pdf>
17. U.S. Department of Energy, "Path to 100% clean electricity," May 2023. <https://www.energy.gov/policy/articles/path-100-clean-electricity>
18. Telefonica, "Telefónica drives energy consumption optimisation through solutions based on artificial intelligence and machine learning," Feb. 2022, https://www.telefonica.com/en/com_munication-room/press-room/telefonica-drives-energy-consumption-optimisation-through-solutions-based-on-artificial-intelligence-and-machine-learning/
19. Ericsson, "Ericsson and FarEasTone: Enhancing 5G with cloud-native solutions." <https://www.ericsson.com/en/cases/2023/ericsson-far-eastone>

20. Ericsson, "Energy Infrastructure Operations." <https://www.ericsson.com/en/managed-services/energy-infrastructure-operations>
21. B. Li, R. Basu Roy, D. Wang, S. Samsi, V. Gadepally, et al., "Toward Sustainable HPC: Carbon Footprint Estimation and Environmental Implications of HPC Systems," Nov. 2023. <https://dl.acm.org/doi/pdf/10.1145/3581784.3607035>
22. A. S. Luccioni, S. Viguier, and A.-L. Ligozat, "Estimating the carbon footprint of BLOOM, a 176B parameter language model," Jan. 2023. <https://www.jmlr.org/papers/volume24/23-0069/23-0069.pdf>
23. 3GPP TS 23.288, "Architecture enhancements for 5G System (5GS) to support network data analytics services," Dec. 2022.
24. 3GPP, "AI/ML Management," Sep. 2023. <https://www.3gpp.org/technologies/ai-ml-management>
25. Qualcomm, "Mobile AI: On-device AI," <https://www.qualcomm.com/products/mobile/snapdragon/smartphones/mobile-ai>
26. ATIS Report, "Advancing Generative AI Implementation in Telecommunications Networks," <https://atis.org/resources/advancing-generative-ai-implementation-in-telecommunications-networks/>
27. Qualcomm, "6G Foundry: How will adaptive intelligence in 6G transform wireless connectivity?" <https://www.qualcomm.com/news/onq/2024/09/6g-foundry-how-will-adaptive-intelligence-in-6g-transform-wireless-connectivity>
28. Next G Alliance Report, "Evolution of Sustainability Indicators for Data Centers and Next Generation Core Networks," https://nextgalliance.org/white_papers/evolution-of-sustainability-indicators-for-data-centers-and-next-generation-core-networks/
29. Next G Alliance Report, "Evolution of Sustainability Indicators for Next Generation Radio Network Technologies," https://nextgalliance.org/white_papers/evolution-of-sustainability-indicators-for-next-generation-radio-network-technologies/
30. Gartner, "Gartner Predicts 30% of Generative AI Projects will be Abandoned After Proof of Concept by End of 2025," Jul. 2024. <https://www.gartner.com/en/newsroom/press-releases/2024-07-29-gartner-predicts-30-percent-of-generative-ai-projects-will-be-abandoned-after-proof-of-concept-by-end-of-2025>
31. Harvard Business Review, "How AI will accelerate the circular economy," Jun. 2023. <https://hbr.org/2023/06/how-ai-will-accelerate-the-circular-economy>
32. E. Strubell, A. Ganesh, and A. McCallum, "Energy and Policy Considerations for Deep Learning in NLP," in Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics, Florence, Italy, Jul. 2019. <https://arxiv.org/abs/1906.02243>
33. Ericsson, "Why AI-powered RAN is an energy efficiency breakthrough," Jan. 2023. <https://www.ericsson.com/en/blog/2023/1/ai-powered-ran-energy-efficiency>
34. ITU-T L.1331, "Assessment of mobile network energy efficiency," <https://www.itu.int/rec/T-REC-L.1331>
35. ETSI ES 203 228, "Assessment of mobile network energy efficiency," https://www.etsi.org/deliver/etsi_es/203200_203299/203228/01.03.01_60/es_203228v010301p.pdf
36. M. Honkala, D. Korpi and J. M. J. Huttunen, "DeepRx: Fully Convolutional Deep Learning Receiver," in IEEE Transactions on Wireless Communications, June 2021.
37. S. Bothe, H. Farooq, J. Forgeat and K. Cyras, "Time-Series Prediction using Nature-Inspired Small Models and Curriculum Learning," 2023. <https://ieeexplore.ieee.org/stamp/stamp.jsp?tp=&number=10294056>
38. OpenAI, "AI and compute," May 2018. <https://openai.com/index/ai-and-compute/>
39. J. B. Jouffray, R. Blasiak, A. V. Norström, H. Österblom, M. Nyström, "The Blue Acceleration: The Trajectory of Human Expansion into the Ocean," Jan. 2020. <https://www.sciencedirect.com/science/article/pii/S2590332219302751>
40. McKinsey & Company, "The growing imperative of energy optimization for telco networks," Feb. 2024. <https://www.mckinsey.com/industries/technology-media-and-telecommunications/our-insights/the-growing-imperative-of-energy-optimization-for-telco-networks>
41. MIT, "New tools are available to help reduce the energy that AI models devour," Oct. 2023. <https://news.mit.edu/2023/new-tools-available-reduce-energy-that-ai-models-devour-1005>
42. Telecoms.com, "NTT Docomo and D-Wave trial give networks some quantum juice," Aug. 2024. <https://www.telecoms.com/wireless-networking/ntt-docomo-and-d-wave-trial-give-networks-some-quantum-juice>

43. B. Al Homssi et al., "Artificial Intelligence Techniques for Next-Generation Massive Satellite Networks" <https://ieeexplore.ieee.org/document/10286330>
44. A. C. Boley and M. Byers, "Satellite mega-constellations create risks in Low Earth Orbit, the atmosphere and on Earth," May 2021. <https://www.nature.com/articles/s41598-021-89909-7>
45. ISO/IEC DTR 20226, "Information technology – Artificial intelligence – Environmental sustainability aspects of AI systems," <https://www.iso.org/standard/86177.html>
46. ITU-T, "Guidelines for Assessing the Impact of Artificial Intelligence on Greenhouse gas emissions". https://www.itu.int/ITU-T/workprog/wp_item.aspx?isn=19564
47. ITU-T, "Deep Learning Computation Energy Efficiency Evaluation Framework and Metrics," https://www.itu.int/ITU-T/workprog/wp_item.aspx?isn=19573
48. ITU-T, "Recommendation for the design of environmentally Sustainable AI-based and XR-based Systems," https://www.itu.int/ITU-T/workprog/wp_item.aspx?isn=19574
49. The White House, "Blueprint for an AI Bill of Rights," <https://www.whitehouse.gov/ostp/ai-bill-of-rights/>
50. Congress.gov, "S.3732 - Artificial Intelligence Environmental Impacts Act of 2024", <https://www.congress.gov/bill/118th-congress/senate-bill/3732/text>

ACKNOWLEDGMENTS

NGA Report Leaders: Gagandeep Bhatti, Nokia
Shruti Bothe, Ericsson

Industry Contributors: Ayodele Damola, Ericsson
Bhushan Joshi, Ericsson
Harsh Malhan, Ericsson
Kamakshi Sridhar, Mavenir
Shohreh Ahvar, Nokia
Chaima Ghribi, Nokia
Olivier Guyot, Nokia
Susanna Kallio, Nokia
Dani Korpi, Nokia
Ibtissam Labriji, Nokia
Alvaro Valcarce, Nokia
Lola Awoniyi-Oteri, Qualcomm
Gavin Horn, Qualcomm
Ammara Essa, Verizon
Jun Shan Wey, Verizon

GREEN G
WORKING GROUP



The mission of the Green G WG is to position North America as the global leader in environmental sustainability by creating a sustainable 6G ecosystem and enabling other industries to reduce greenhouse gases and energy consumption, limit land and water use, and move towards circular economy.

Green G Working Group Leadership

Chair, Bhushan Joshi, Ericsson

Vice Chair, Ralf Bendlin, AT&T


Vice Chair, Gagandeep Bhatti, Nokia

Ian Deakin, ATIS

Green G Working Group Membership

Analog Devices, Inc.
Anritsu
AT&T
ATIS
Bell Canada
CableLabs
Case Western Reserve University
Charter Communications
Ciena Corporation
Cisco
Cohere Technologies
DeepSig
DOCOMO Innovations Inc.
Ericsson
Florida Atlantic University
Futurewei
George Mason University
Georgetown University
Google
Hewlett Packard Enterprise
IMEC
Intel Corporation
InterDigital Communications Corporation
Iowa State University
ITRI
Johns Hopkins University Applied
Physics Laboratory
Keysight Technologies, Inc
LG Electronics
Lockheed Martin
Mavenir
Mediatek Inc
Microsoft Corporation
MITRE Corporation
Motorola Mobility LLC (A Lenovo
Company)
Murata Manufacturing Co., Ltd.
NIST
Nokia
North Carolina State University
Northwestern University
NTIA
Ofinno
Old Dominion University
Purdue University
Qualcomm Incorporated
Samsung Research America
Schneider Electric
Sharp Labs of America
Spreadtrum Communication USA, Inc.
T-Mobile USA
TDS
TELUS
University of Manitoba
University of Notre Dame
Verizon
Verizon Wireless
Viavi Solutions
Virginia Tech
VT-ARC
WIN Semiconductors Corp.
Wireless Internet of Things at
Northeastern University

NEXT G ALLIANCE REPORTS



Next G Alliance Report: 6G Technologies



Next G Alliance Report: 6G Applications and Use Cases



Next G Alliance Report: Roadmap to 6G



Green G: The Path Toward Sustainable 6G



6G Market Development: A North American Perspective



Next G Alliance Report: AI-Native Wireless Networks



Next G Alliance Report: Sustainable 6G Connectivity – A Powerful Means of Doing Good



Next G Alliance Report: 6G Distributed Cloud and Communications System



Next G Alliance Report: Trust, Security, and Resilience for 6G Systems



Next G Alliance Report: Digital World Experiences



Next G Alliance Report: Cost-Efficient Solutions



Next G Alliance Report: 6G Radio Technology Part I: Basic Radio Technologies




Next G Alliance Report: 6G Radio Technology Part II: Basic Radio Technologies




Next G Alliance Report: 6G Technologies for Wide-Area Cloud Evolution



6G Market Management and Orchestration




Next G Alliance Report: Beyond Speed: Promoting Social and Economic Opportunities through 6G and Beyond



Next G Alliance Report: 6G Roadmap for Vertical Industries



Next G Alliance Report: Channel Measurements and Modeling for Joint/Integrated Communication and Sensing, as well as 7-24 GHz Communication



Next G Alliance Report: 6G Sustainability KPI North American 6G Roadmap Priorities



Next G Alliance Report: Distributed Sensing and Communications




Next G Alliance Report: Network-Enabled Robotic and Autonomous Systems



Next G Alliance Report: Multi-Sensory Extended Reality (XR) in 6G



Next G Alliance Report: Shaping Tomorrow: The Evolution of Personalized Digital Experiences Through 6G Technologies



Next G Alliance Report: Video Codes for 6G Machine-Type Communications: Improving the Quality of Critical Application Roles



Next G Alliance Report: Spectrum Working Group: Spectrum Access Mechanisms



Next G Alliance Report: Spectrum Needs for 6G



Next G Alliance Report: 6G Spectrum Considerations



Next G Alliance Report: Terminology for Frequency Ranges



Next G Alliance Report: Evolution of Sustainability Indicators for Data Centers and Next Generation Core Networks



Next G Alliance Report: Evolution of Sustainability Indicators for Next-Generation Radio Network Technologies



Next G Alliance Report: 6G Sustainability KPI Assessment Introduction and Gap Analysis

COPYRIGHT
AND
DISCLAIMER

Published February 2025

Copyright © 2025 by Alliance for Telecommunications Industry Solutions

All rights reserved.

Alliance for Telecommunications Industry Solutions

1200 G Street, NW, Suite 500

Washington, DC 20005

No part of this publication may be reproduced in any form, in an electronic retrieval system or otherwise, without the prior written permission of the publisher. For information, contact ATIS at (202) 628-6380. ATIS is online at <http://www.atis.org>.

The information provided in this document is directed solely to professionals who have the appropriate degree of experience to understand and interpret its contents in accordance with generally accepted engineering or other professional standards and applicable regulations. No recommendation as to products or vendors is made or should be implied.

NO REPRESENTATION OR WARRANTY IS MADE THAT THE INFORMATION IS TECHNICALLY ACCURATE OR SUFFICIENT OR CONFORMS TO ANY STATUTE, GOVERNMENTAL RULE OR REGULATION, AND FURTHER, NO REPRESENTATION OR WARRANTY IS MADE OF MERCHANTABILITY OR FITNESS FOR ANY PARTICULAR PURPOSE OR AGAINST INFRINGEMENT OF INTELLECTUAL PROPERTY RIGHTS. ATIS SHALL NOT BE LIABLE, BEYOND THE AMOUNT OF ANY SUM RECEIVED IN PAYMENT BY ATIS FOR THIS DOCUMENT, AND IN NO EVENT SHALL ATIS BE LIABLE FOR LOST PROFITS OR OTHER INCIDENTAL OR CONSEQUENTIAL DAMAGES. ATIS EXPRESSLY ADVISES THAT ANY AND ALL USE OF OR RELIANCE UPON THE INFORMATION PROVIDED IN THIS DOCUMENT IS AT THE RISK OF THE USER.

NOTE - The user's attention is called to the possibility that compliance with this document may require use of an invention covered by patent rights. By publication of this document, no position is taken with respect to whether use of an invention covered by patent rights will be required, and if any such use is required no position is taken regarding the validity of this claim or any patent rights in connection therewith. Please refer to www.atis.org/legal/patentinfo.asp to determine if any statement has been filed by a patent holder indicating a willingness to grant a license either without compensation or on reasonable and non-discriminatory terms and conditions to applicants desiring to obtain a license.

Building the foundation for
North American leadership in
6G and beyond.

nextgalliance.org

