Federated Deep Reinforcement Learning for Adaptive Spectrum Allocation in Dense 5G-A Wireless Networks

Zaid J. Reed
Department of Electrical and Computer Engineering, University of Nebraska–Omaha
zaidreed@unomaha.edu

Kriti Grover
School of Computing and Information Sciences, University of North Texas
kriti-grover87@unt.edu

Abstract

The evolution of fifth-generation advanced wireless systems has intensified the complexity of spectrum allocation in dense heterogeneous communication environments characterized by massive device connectivity, ultra-low latency requirements, network slicing, edge intelligence, and highly dynamic traffic behavior. Traditional spectrum management approaches based on static assignment, centralized optimization, or isolated learning architectures have demonstrated limited scalability under rapidly changing radio conditions and geographically distributed service demands. This paper investigates the integration of federated deep reinforcement learning as a distributed intelligence paradigm for adaptive spectrum allocation in dense 5G-A wireless networks. The study explores how federated collaborative learning mechanisms enable decentralized base stations, edge nodes, and network controllers to jointly optimize spectrum utilization while preserving local operational autonomy and minimizing raw data exchange. The paper develops a system-level analytical framework examining the interaction between spectrum scarcity, edge orchestration, interference management, mobility-aware optimization, and trust-aware policy coordination across heterogeneous radio access infrastructures. Particular attention is given to governance challenges associated with fairness, privacy preservation, model drift, communication overhead, and energy sustainability within large-scale wireless ecosystems. The analysis further evaluates the architectural implications of integrating deep reinforcement learning with federated coordination under high-density urban deployments, industrial communication systems, vehicular networking environments, and critical infrastructure services. Beyond performance optimization, the paper emphasizes the broader socio-technical implications of intelligent spectrum orchestration, including regulatory adaptation, infrastructure inequality, environmental efficiency, and operational resilience. The findings indicate that federated deep reinforcement learning provides a viable pathway toward scalable, adaptive, and resilient spectrum governance for next-generation wireless ecosystems, particularly when supported by cross-layer coordination, policy-aware deployment strategies, and robust edge intelligence frameworks.

Keywords

Federated learning, deep reinforcement learning, adaptive spectrum allocation, 5G-Advanced, dense wireless networks, edge intelligence, spectrum management, network orchestration, distributed AI, wireless infrastructure resilience.

1. Introduction

The rapid expansion of wireless communication infrastructures has transformed spectrum allocation into one of the most critical operational challenges in modern digital ecosystems. Fifth-generation advanced wireless systems, commonly referred to as 5G-A networks, have introduced unprecedented levels of device density, mobility heterogeneity, ultra-reliable low-latency communication requirements, and edge-native service architectures. These developments have significantly increased the complexity of radio resource coordination across urban metropolitan environments, industrial automation systems, intelligent transportation infrastructures, and large-scale Internet of Things deployments. Unlike earlier wireless generations that primarily emphasized throughput enhancement and broad connectivity expansion, contemporary 5G-A systems must simultaneously support latency-sensitive applications, autonomous systems coordination, immersive communication platforms, and mission-critical industrial operations within increasingly congested radio environments [1][2].

Traditional spectrum allocation methods have historically relied on centralized optimization frameworks, fixed allocation strategies, heuristic scheduling mechanisms, or static interference mitigation models. While these approaches performed adequately in relatively stable communication environments, they encounter severe limitations in dense and highly dynamic wireless ecosystems characterized by rapidly fluctuating traffic loads, distributed edge intelligence, and heterogeneous service priorities [3]. The growth of edge computing and decentralized network architectures has further complicated centralized spectrum governance due to communication overhead, privacy constraints, and the operational inefficiencies associated with transmitting large-scale radio environment data to centralized cloud infrastructures [4].

Artificial intelligence has emerged as a transformative paradigm for adaptive wireless network management. In particular, deep reinforcement learning has attracted substantial attention because of its capability to autonomously optimize decision-making processes in uncertain and dynamically evolving environments [5]. By continuously interacting with radio environments, reinforcement learning agents can adapt transmission policies, allocate communication channels, and optimize spectrum utilization without requiring explicit analytical modeling of complex wireless conditions. However, standalone deep reinforcement learning architectures often face practical deployment limitations in large-scale wireless systems because centralized training introduces scalability bottlenecks, privacy concerns, and excessive communication latency [6].

Federated learning has consequently emerged as a promising distributed intelligence framework capable of enabling collaborative model optimization without direct exchange of raw operational data [7]. Under federated learning paradigms, geographically distributed nodes independently train local models using locally observed network conditions while periodically sharing model parameters or gradients with aggregation servers. This distributed learning mechanism offers substantial advantages for wireless infrastructures where privacy preservation, low-latency decision-making, and operational autonomy are essential. The integration of federated learning with deep reinforcement learning therefore represents a potentially transformative approach for adaptive spectrum allocation in dense 5G-A ecosystems [8].

The relevance of federated deep reinforcement learning extends beyond technical optimization alone. Spectrum allocation increasingly intersects with broader socio-technical concerns involving infrastructure fairness, regulatory compliance, sustainability, energy consumption, digital inequality, and governance transparency. Intelligent spectrum coordination mechanisms influence not only network efficiency but also economic accessibility, emergency response capabilities, industrial productivity, and societal resilience under large-scale digital dependence [9]. Consequently, the deployment of federated intelligence architectures within wireless infrastructures must be examined through multidimensional analytical perspectives encompassing technical, organizational, and policy-oriented considerations.

This paper develops a comprehensive systems-oriented investigation of federated deep reinforcement learning for adaptive spectrum allocation in dense 5G-A wireless networks. The analysis emphasizes architectural trade-offs, distributed intelligence coordination, deployment challenges, operational resilience, and long-term infrastructure implications. Rather than focusing exclusively on algorithmic performance metrics, the paper situates federated reinforcement learning within the broader context of evolving wireless ecosystems characterized by edge-native intelligence, infrastructure decentralization, and adaptive communication governance.

The remainder of this paper is organized as follows. Section 2 reviews the evolution of spectrum allocation mechanisms and the emergence of intelligent wireless resource orchestration. Section 3 examines the architectural foundations of federated deep reinforcement learning within dense 5G-A environments. Section 4 analyzes adaptive spectrum allocation mechanisms under distributed learning conditions. Section 5 discusses scalability, fairness, and robustness challenges. Section 6 explores energy sustainability and environmental implications. Section 7 investigates regulatory and governance dimensions. Section 8 presents deployment-oriented case analyses across heterogeneous wireless domains. Section 9 discusses future research trajectories and emerging infrastructure paradigms. Finally, Section 10 concludes the paper by synthesizing the broader implications of federated deep reinforcement learning for next-generation wireless ecosystems.

2. Evolution of Spectrum Allocation and Intelligent Wireless Resource Management

Spectrum allocation has historically represented a foundational challenge in wireless communication engineering because radio frequency resources are inherently limited while communication demand continuously expands. Earlier wireless generations largely relied on rigid allocation frameworks governed by fixed spectrum licensing models and centrally managed radio planning strategies. These approaches prioritized stability and predictability over adaptability, reflecting the relatively static traffic characteristics and limited device diversity of legacy communication systems [10].

As wireless ecosystems evolved toward broadband multimedia communication and ubiquitous mobile connectivity, dynamic spectrum management gradually became necessary. Fourth-generation wireless systems introduced increasingly sophisticated scheduling techniques designed to improve throughput efficiency and mitigate interference across heterogeneous radio access networks. Nevertheless, these systems still largely depended on predefined optimization heuristics and rule-based coordination frameworks that struggled to adapt effectively under rapidly fluctuating network conditions [11].

The transition toward 5G and subsequently 5G-A infrastructures fundamentally altered the operational landscape of wireless communication systems. Contemporary wireless ecosystems support highly heterogeneous applications ranging from autonomous transportation and industrial robotics to immersive augmented reality and large-scale environmental sensing. These applications exhibit vastly different latency sensitivities, reliability requirements, mobility patterns, and bandwidth consumption characteristics [12]. Consequently, spectrum allocation mechanisms must continuously adapt to multidimensional operational conditions rather than optimizing for singular performance metrics such as throughput or spectral efficiency alone.

The densification of wireless infrastructures has further intensified resource coordination challenges. Urban communication environments now contain overlapping layers of macro cells, small cells, Wi-Fi networks, edge computing nodes, private industrial communication systems, and vehicular communication infrastructures. The resulting radio environments exhibit substantial temporal variability and interference complexity, rendering conventional optimization approaches increasingly inadequate [13].

Machine learning consequently emerged as a strategic tool for enabling adaptive wireless resource management. Supervised learning methods initially supported predictive traffic modeling and anomaly detection applications, while unsupervised learning facilitated network clustering and behavior analysis. However, reinforcement learning gained particular relevance because wireless environments are fundamentally dynamic and sequential decision-oriented systems in which agents must continuously adapt policies based on environmental feedback [14].

Deep reinforcement learning extended the capabilities of conventional reinforcement learning by incorporating deep neural network architectures capable of processing high-dimensional state spaces. This advancement enabled wireless agents to interpret complex radio environment conditions, traffic dynamics, and interference relationships with significantly improved adaptability [15]. Deep reinforcement learning models demonstrated promising capabilities in channel selection, power allocation, handover optimization, and network slicing management within simulated wireless environments.

Despite these advancements, centralized deep reinforcement learning architectures introduced substantial practical concerns. Centralized training frameworks require large-scale data aggregation across geographically distributed infrastructures, generating communication overhead and latency accumulation that can undermine real-time responsiveness [16]. Furthermore, centralized data collection raises privacy and security concerns because wireless operational data may reveal sensitive information regarding user mobility patterns, industrial processes, or critical infrastructure activities.

Federated learning emerged in response to these limitations by enabling distributed model training without centralized raw data aggregation [7]. Initially developed for privacy-sensitive mobile computing applications, federated learning rapidly gained relevance within wireless communication systems because distributed network infrastructures naturally align with collaborative learning architectures. Under federated learning paradigms, local nodes train models independently using locally observed conditions while sharing only model updates with coordination servers.

The convergence of federated learning and deep reinforcement learning introduced new possibilities for distributed wireless intelligence. Federated deep reinforcement learning enables geographically distributed wireless agents to collaboratively improve resource management policies while preserving localized operational autonomy [17]. This paradigm aligns particularly well with 5G-A environments characterized by edge-native computation, decentralized network slicing, and infrastructure heterogeneity.

Recent research has demonstrated growing interest in applying federated deep reinforcement learning to spectrum allocation, interference mitigation, edge caching, vehicular communication coordination, and mobility-aware network optimization [18]. However, many existing studies primarily focus on algorithmic performance evaluation under constrained simulation settings. Broader system-level analyses examining governance implications, deployment scalability, fairness considerations, and socio-technical infrastructure impacts remain comparatively underdeveloped.

The evolution of intelligent spectrum management therefore reflects a broader transition from static engineering optimization toward adaptive, distributed, and collaborative infrastructure intelligence. Contemporary wireless ecosystems increasingly resemble complex socio-technical systems requiring coordinated interaction among technical architectures, regulatory institutions, operational stakeholders, and evolving service demands. Federated deep reinforcement learning represents not merely an optimization mechanism but a foundational shift in how wireless infrastructures govern shared communication resources under conditions of extreme complexity and uncertainty.

3. Architectural Foundations of Federated Deep Reinforcement Learning in 5G-A Systems

The integration of federated deep reinforcement learning into dense 5G-A wireless networks requires a multilayered architectural framework capable of coordinating distributed intelligence across heterogeneous communication infrastructures. Unlike conventional centralized wireless optimization systems, federated deep reinforcement learning architectures distribute computational responsibility across edge nodes, radio access networks, local controllers, and aggregation entities while maintaining collaborative policy adaptation mechanisms [19].

At the infrastructural level, 5G-A ecosystems are characterized by highly distributed network topologies consisting of macro base stations, small cells, edge computing clusters, intelligent reflecting surfaces, network slicing controllers, and software-defined networking components. These heterogeneous entities collectively participate in radio resource management processes while operating under varying computational capacities, latency constraints, and mobility conditions [20]. Federated learning frameworks align naturally with this distributed topology because local spectrum management decisions can be optimized directly at edge locations where operational conditions are observed in real time.

Deep reinforcement learning agents embedded within local base stations or edge controllers continuously interact with radio environments to optimize spectrum allocation decisions. These agents observe local traffic patterns, interference conditions, channel occupancy states, user mobility dynamics, and service quality indicators. Based on these observations, local agents iteratively refine allocation strategies designed to maximize long-term network performance objectives while minimizing congestion and interference accumulation [21].

Federated coordination mechanisms enable these distributed agents to exchange learned intelligence without transmitting sensitive operational datasets. Instead of sharing raw radio measurements or user-level information, local agents periodically transmit model updates to federated aggregation layers responsible for synthesizing global policy representations [22]. Aggregated knowledge is subsequently redistributed to participating nodes, enabling collaborative adaptation while preserving localized operational privacy.

This distributed architecture offers several strategic advantages within dense wireless ecosystems. First, local decision-making reduces latency accumulation associated with centralized cloud processing. Spectrum allocation decisions often require near-real-time responsiveness because radio conditions fluctuate rapidly under mobility-intensive environments. Edge-native federated reinforcement learning therefore improves responsiveness by maintaining intelligence near communication endpoints [23].

Second, federated architectures improve scalability under large-scale device densities. Contemporary urban wireless environments may involve millions of simultaneously connected devices generating continuously evolving traffic patterns. Centralized learning frameworks face severe bottlenecks when aggregating such large-scale operational data streams. Federated learning mitigates these constraints by distributing training workloads across local infrastructures [24].

Third, distributed intelligence improves resilience against partial infrastructure failures. Centralized spectrum management systems represent potential single points of failure under cyberattacks, network disruptions, or infrastructure outages. Federated architectures distribute operational intelligence across multiple nodes, thereby improving continuity during localized failures [25].

Nevertheless, federated deep reinforcement learning introduces substantial architectural complexity. Model synchronization across heterogeneous wireless infrastructures remains challenging because local environments often exhibit non-independent and non-identically distributed data characteristics. Urban commercial districts, industrial zones, transportation corridors, and residential areas may experience radically different traffic dynamics and interference behaviors [26]. Consequently, globally aggregated models may struggle to generalize effectively across highly heterogeneous operational conditions.

Communication overhead also emerges as a critical challenge. Although federated learning reduces raw data transmission requirements, periodic model synchronization still consumes network resources. Under extremely dense deployments involving thousands of participating nodes, federated communication traffic may itself contribute to congestion accumulation [27].

Efficient synchronization scheduling and hierarchical aggregation strategies therefore become essential for maintaining operational feasibility.

Edge computational heterogeneity further complicates deployment. Some wireless nodes possess advanced hardware acceleration capabilities, while others operate under constrained energy and processing budgets. Federated reinforcement learning architectures must therefore accommodate asymmetric training capacities and dynamically adapt participation mechanisms according to local infrastructure conditions [28].

Security and trust management also represent critical architectural considerations. Federated learning systems may be vulnerable to model poisoning attacks, adversarial manipulation, or malicious gradient injection. In wireless infrastructures supporting critical societal services, compromised learning coordination mechanisms could destabilize communication reliability or introduce discriminatory allocation behaviors [29]. Consequently, secure aggregation protocols, trust-aware validation mechanisms, and anomaly detection frameworks become integral components of resilient federated architectures.

Cross-layer integration further influences architectural effectiveness. Spectrum allocation decisions interact with higher-layer network slicing policies, edge application orchestration mechanisms, transport protocols, and service-level agreements. Federated reinforcement learning architectures therefore require coordination across multiple infrastructural layers rather than operating solely at the physical or medium access control levels [30].

The architectural significance of federated deep reinforcement learning extends beyond technical optimization alone. These systems effectively redistribute decision-making authority across wireless ecosystems, shifting intelligence from centralized operators toward collaborative edge infrastructures. This transition raises broader governance questions concerning accountability, transparency, interoperability, and infrastructure sovereignty within increasingly autonomous communication systems.

4. Adaptive Spectrum Allocation Under Distributed Learning Conditions

Adaptive spectrum allocation within dense 5G-A environments involves continuously balancing competing operational objectives across highly dynamic communication ecosystems. Federated deep reinforcement learning introduces a distributed intelligence framework capable of coordinating these objectives under conditions of environmental uncertainty, infrastructure heterogeneity, and fluctuating service demand [31].

Dense wireless environments exhibit substantial temporal variability due to user mobility, application diversity, environmental interference, and traffic burstiness. Traditional static allocation mechanisms struggle because spectrum utilization patterns can change dramatically within short temporal intervals. Federated reinforcement learning enables continuous adaptation by allowing local agents to refine allocation strategies based on evolving environmental observations [32].

Local reinforcement learning agents typically optimize multidimensional performance objectives including throughput stability, latency minimization, fairness preservation, interference mitigation, and energy efficiency. Unlike rule-based scheduling systems, reinforcement learning agents gradually discover allocation policies through repeated environmental interaction rather than relying on predefined analytical assumptions [33]. This adaptive capability is particularly valuable within complex urban wireless ecosystems where radio conditions are influenced by unpredictable environmental factors such as physical obstructions, mobility congestion, and cross-network interference.

Federated coordination enhances this adaptability by enabling geographically distributed agents to share learned operational intelligence. For example, spectrum congestion patterns identified in one urban district may provide valuable insights for neighboring districts experiencing similar mobility behaviors or traffic surges. Collaborative learning therefore accelerates policy convergence and improves generalization across heterogeneous operational environments [34].

Network slicing further complicates adaptive spectrum allocation because different service categories exhibit divergent performance requirements. Enhanced mobile broadband applications prioritize high throughput, while ultra-reliable low-latency communication services require deterministic responsiveness. Massive machine-type communication environments prioritize scalability and energy efficiency. Federated reinforcement learning architectures can dynamically adapt allocation policies according to evolving slice-specific conditions while maintaining collaborative coordination across shared infrastructures [35].

Interference management represents another critical operational challenge. Dense small-cell deployments significantly increase spectral reuse opportunities but simultaneously intensify interference complexity. Federated reinforcement learning agents can collaboratively identify interference patterns and optimize channel allocation strategies based on distributed environmental observations [36]. Unlike isolated local optimization frameworks, federated coordination enables broader spatial awareness across overlapping communication domains.

Mobility-aware spectrum allocation is increasingly important within vehicular communication systems, drone-enabled networks, and mobile edge computing environments. User mobility continuously alters channel quality, interference relationships, and traffic distribution patterns. Federated reinforcement learning facilitates adaptive handover optimization and mobility-sensitive spectrum coordination by enabling local agents to learn from geographically distributed mobility experiences [37].

The integration of edge computing substantially enhances the feasibility of adaptive distributed learning. Edge infrastructures provide localized computational resources capable of supporting real-time reinforcement learning operations without excessive dependence on centralized cloud systems. This architectural proximity reduces latency accumulation while improving responsiveness under rapidly changing radio conditions [38].

However, adaptive spectrum allocation through federated reinforcement learning also introduces significant operational trade-offs. Continuous adaptation may increase policy instability under highly volatile environments, particularly when local agents encounter conflicting optimization incentives. Excessive exploration behaviors may temporarily degrade service quality during learning phases [39]. Balancing exploration and operational stability therefore becomes essential within production-grade wireless infrastructures.

Fairness preservation further complicates allocation dynamics. Reinforcement learning systems optimized primarily for aggregate throughput may inadvertently disadvantage low-density regions, low-priority users, or resource-constrained devices. Federated architectures may amplify such disparities if dominant nodes disproportionately influence global model aggregation [40]. Fairness-aware coordination mechanisms are therefore necessary to prevent systematic inequities within adaptive spectrum governance.

Communication synchronization frequency also influences operational efficiency. Frequent federated updates improve collaborative adaptation but increase network overhead. Infrequent synchronization reduces communication costs but may slow policy convergence under rapidly evolving conditions [41]. Dynamic synchronization scheduling strategies are consequently required to balance collaborative intelligence sharing against communication resource consumption.

Privacy preservation introduces additional complexity. Although federated learning minimizes raw data exchange, model updates may still leak sensitive operational information through inference attacks or gradient reconstruction techniques [42]. Privacy-enhancing mechanisms such as differential privacy, secure multiparty computation, and encrypted aggregation protocols must therefore be integrated carefully without excessively degrading learning effectiveness.

The adaptive allocation capabilities enabled by federated deep reinforcement learning ultimately reflect a broader transformation in wireless infrastructure management. Contemporary communication systems increasingly operate as self-organizing intelligent ecosystems capable of collaborative environmental adaptation rather than centrally programmed operational control. This shift fundamentally redefines how spectrum resources are governed within digitally interconnected societies.

5. Scalability, Robustness, and Fairness Challenges

The deployment of federated deep reinforcement learning within dense 5G-A wireless ecosystems introduces substantial scalability, robustness, and fairness challenges that extend far beyond algorithmic optimization. While distributed intelligence frameworks offer promising adaptability and resilience advantages, their practical implementation across large-scale communication infrastructures requires careful consideration of operational constraints, systemic inequalities, and infrastructural vulnerabilities [43].

Scalability represents one of the most immediate concerns in federated wireless learning environments. Contemporary 5G-A systems may involve thousands of distributed nodes participating in collaborative learning processes across metropolitan regions, industrial campuses, transportation corridors, and critical infrastructure zones. Coordinating model synchronization among such large-scale distributed agents can generate substantial communication overhead and computational complexity [44].

The heterogeneity of participating infrastructures further complicates scalability management. Wireless nodes vary considerably in processing capabilities, energy availability, storage capacity, and connectivity stability. Some edge devices may support advanced hardware acceleration for deep learning operations, while others operate under constrained embedded system environments. Federated reinforcement learning architectures must therefore accommodate asynchronous participation and variable training capacities without destabilizing collaborative policy convergence [45].

Model convergence itself becomes increasingly difficult under non-independent and non-identically distributed operational conditions. Different geographic regions exhibit distinct traffic behaviors, mobility patterns, environmental interference characteristics, and service demand distributions. Local reinforcement learning agents may therefore optimize divergent policies that are difficult to reconcile through global aggregation mechanisms [46]. Excessive model divergence can slow convergence rates, reduce generalization effectiveness, and destabilize collaborative coordination processes.

Communication latency also influences scalability. Dense wireless ecosystems already experience substantial signaling overhead associated with mobility management, network slicing coordination, and edge orchestration processes. Federated synchronization introduces additional signaling requirements that may compete with primary communication services for limited network resources [47]. Hierarchical federated architectures and adaptive aggregation strategies consequently become essential for maintaining operational scalability.

Robustness challenges emerge from the highly dynamic and adversarial nature of wireless environments. Communication infrastructures supporting transportation systems, industrial automation, healthcare services, and emergency response operations require exceptional operational reliability. Federated reinforcement learning systems must therefore maintain stable performance despite environmental uncertainty, infrastructure failures, cyber threats, and unpredictable traffic fluctuations [48].

Model poisoning attacks represent a particularly serious concern. Malicious participants may intentionally manipulate local model updates to degrade global policy effectiveness or introduce harmful allocation behaviors. In critical communication infrastructures, compromised spectrum allocation policies could disrupt public safety services, destabilize industrial operations, or create systemic communication failures [49]. Trust-aware aggregation mechanisms and anomaly detection frameworks are therefore essential components of secure federated wireless intelligence.

Adversarial interference environments further complicate robustness management. Wireless infrastructures are vulnerable to jamming attacks, spoofing behaviors, and malicious interference generation capable of distorting environmental observations used by reinforcement learning agents. Federated architectures may inadvertently propagate corrupted learning signals across distributed infrastructures if validation mechanisms are insufficiently robust [50].

Environmental uncertainty also introduces operational instability. Reinforcement learning systems rely heavily on exploratory behavior during policy optimization, yet excessive exploration may temporarily degrade service quality in production environments. Balancing adaptive learning against operational stability becomes particularly challenging in mission-critical communication systems where reliability requirements are stringent [51].

Fairness constitutes another foundational challenge within federated spectrum allocation systems. Communication infrastructures increasingly influence economic participation, educational access, healthcare delivery, and social inclusion. Consequently, spectrum allocation policies carry significant societal implications beyond technical efficiency metrics [52].

Reinforcement learning systems optimized primarily for aggregate network performance may unintentionally reinforce infrastructural inequalities. High-density commercial regions generating substantial traffic demand may receive disproportionate spectrum prioritization, while rural communities, low-income regions, or lower-priority services experience degraded performance [53]. Such disparities could exacerbate existing digital inequalities within increasingly network-dependent societies.

Federated learning architectures may also reproduce asymmetrical influence structures during model aggregation. Large network operators possessing extensive infrastructure resources may dominate collaborative learning processes, thereby marginalizing smaller providers or localized communication systems [54]. Governance frameworks ensuring equitable participation and balanced influence distribution are therefore necessary for maintaining fairness within collaborative wireless intelligence ecosystems.

Bias accumulation further complicates fairness preservation. Reinforcement learning agents trained predominantly on specific operational environments may generalize poorly across underrepresented user populations or infrastructure conditions. Urban-centric training data, for instance, may inadequately capture rural communication dynamics or industrial operational requirements [55]. Diverse participation mechanisms and fairness-aware optimization objectives therefore become critical for preventing discriminatory allocation behaviors.

Energy sustainability introduces additional systemic constraints. Federated deep reinforcement learning requires substantial computational resources for continuous local training, model synchronization, and distributed inference operations. Large-scale deployment across dense wireless infrastructures may significantly increase energy consumption and carbon emissions if computational efficiency is not prioritized [56]. Sustainable infrastructure design therefore becomes inseparable from scalable federated intelligence deployment.

These interconnected challenges highlight that federated deep reinforcement learning should not be viewed merely as a technical optimization framework. Rather, it represents a complex socio-technical infrastructure transformation requiring coordinated consideration of scalability engineering, resilience governance, fairness regulation, sustainability management, and institutional accountability.

6. Energy Sustainability and Environmental Implications

The environmental sustainability of intelligent wireless infrastructures has emerged as a central concern in the evolution of 5G-A communication systems. Dense wireless deployments require substantial computational and communication resources, while federated deep reinforcement learning introduces additional energy demands associated with distributed model training, synchronization, and inference operations [57]. Consequently, adaptive spectrum allocation architectures must increasingly balance performance optimization against environmental responsibility and long-term infrastructure sustainability.

The expansion of edge-native intelligence significantly alters the energy dynamics of wireless ecosystems. Traditional centralized cloud architectures concentrate computational workloads within hyperscale data centers, whereas federated learning distributes training responsibilities across geographically dispersed edge infrastructures. While distributed computation may reduce backbone communication overhead, it also increases aggregate energy consumption across local devices and edge nodes [58].

Deep reinforcement learning itself is computationally intensive because policy optimization involves continuous environmental interaction, neural network updates, and long-term reward estimation processes. When multiplied across thousands of participating wireless nodes, cumulative computational energy consumption can become substantial [59]. This challenge is particularly pronounced in dense urban environments where large numbers of small cells and edge devices operate continuously under high traffic loads.

Adaptive spectrum allocation can nevertheless contribute positively to sustainability objectives when designed appropriately. Intelligent allocation policies may reduce unnecessary transmission power, minimize interference-related retransmissions, optimize sleep scheduling mechanisms, and improve spectral efficiency. These optimizations can collectively reduce overall network energy consumption despite the computational costs associated with distributed learning [60].

Federated coordination further offers opportunities for environmentally efficient intelligence sharing. By avoiding large-scale centralized raw data aggregation, federated learning reduces backbone communication traffic and associated energy expenditures. Localized processing also enables context-sensitive optimization strategies tailored to regional infrastructure conditions, potentially improving overall operational efficiency [61].

Energy-aware reinforcement learning has consequently emerged as an important research direction within intelligent wireless systems. Rather than optimizing solely for throughput or latency, reinforcement learning agents increasingly incorporate energy efficiency objectives into policy adaptation processes. Such multidimensional optimization frameworks aim to balance communication quality against environmental sustainability under varying operational conditions [62].

Renewable energy integration further influences sustainability considerations. Many contemporary edge infrastructures increasingly incorporate solar, wind, or hybrid renewable power systems. Federated reinforcement learning agents may adapt spectrum allocation strategies according to fluctuating renewable energy availability, thereby improving environmental efficiency and reducing dependence on carbon-intensive energy sources [63].

However, sustainability trade-offs remain significant. More sophisticated learning architectures often require larger neural networks and more frequent synchronization operations, increasing computational intensity. Similarly, achieving extremely low-latency communication performance may necessitate continuous infrastructure activation and aggressive resource provisioning, thereby increasing energy demand [64].

Hardware lifecycle considerations also affect environmental sustainability. The deployment of advanced edge accelerators, specialized AI processors, and dense communication hardware contributes to electronic waste generation and resource extraction pressures. Rapid technological obsolescence within wireless industries may further intensify environmental burdens if infrastructure upgrade cycles remain excessively short [65].

Federated learning may nevertheless extend hardware longevity by enabling localized optimization without requiring continuous centralized infrastructure expansion. Distributed intelligence frameworks can improve the utilization efficiency of existing communication assets, potentially reducing the need for aggressive infrastructure overprovisioning [66].

Environmental sustainability also intersects with broader societal resilience concerns. Climate-related disruptions increasingly threaten communication infrastructures through extreme weather events, heat stress, flooding, and energy instability. Adaptive federated learning systems capable of decentralized coordination may improve infrastructure resilience under environmentally unstable conditions [67]. Distributed intelligence reduces dependence on centralized coordination facilities that may themselves become vulnerable during climate-related disruptions.

The relationship between sustainability and fairness is equally important. Energy-efficient spectrum allocation strategies must avoid disproportionately degrading communication quality for marginalized regions or underserved populations. Sustainable optimization should therefore incorporate equity-oriented objectives rather than focusing exclusively on aggregate energy reduction metrics [68].

Policy and regulatory frameworks increasingly influence sustainable wireless infrastructure development. Governments and international standards organizations are introducing environmental performance expectations for communication systems, including energy reporting requirements, carbon reduction targets, and infrastructure efficiency standards [69]. Federated deep reinforcement learning systems will consequently need to operate within evolving regulatory environments emphasizing environmental accountability alongside communication performance.

The environmental implications of intelligent wireless infrastructures ultimately extend beyond operational energy consumption alone. They encompass broader questions concerning technological sustainability, resource allocation ethics, digital infrastructure resilience, and the long-term environmental consequences of increasingly autonomous communication ecosystems. Adaptive spectrum allocation must therefore be evaluated not merely through technical efficiency metrics but through holistic sustainability-oriented infrastructure perspectives.

7. Governance, Regulation, and Policy Implications

The deployment of federated deep reinforcement learning within dense 5G-A wireless ecosystems raises profound governance and regulatory questions concerning accountability, transparency, sovereignty, and institutional control over increasingly autonomous communication infrastructures. Spectrum allocation has historically been governed through centralized regulatory institutions and licensed operational frameworks. However, distributed intelligent allocation mechanisms fundamentally alter how spectrum governance is operationalized within digitally interconnected societies [70].

One of the central governance challenges involves accountability in autonomous decision-making systems. Federated reinforcement learning agents continuously adapt spectrum allocation policies based on evolving environmental observations and collaborative model updates. While this adaptability improves operational efficiency, it may also reduce transparency regarding how specific allocation decisions are generated [71]. In critical communication environments, opaque decision-making processes may complicate regulatory oversight and operational auditing.

Regulators increasingly require explainability and traceability within AI-enabled infrastructures. Communication systems supporting emergency response operations, healthcare services, transportation coordination, and industrial automation cannot rely entirely on black-box optimization mechanisms without sufficient governance safeguards [72]. Federated learning architectures therefore require explainable policy coordination frameworks capable of supporting operational transparency and institutional accountability.

Data sovereignty constitutes another critical policy concern. Wireless operational data often reflects sensitive information regarding mobility patterns, industrial activities, and infrastructure utilization. Federated learning partially mitigates privacy concerns by avoiding centralized raw data aggregation. Nevertheless, collaborative model updates may still reveal sensitive operational characteristics through inference attacks or adversarial reconstruction techniques [73]. Regulatory frameworks governing privacy protection and cross-border data exchange consequently influence the feasibility of large-scale federated learning deployment.

Geopolitical considerations further complicate governance dynamics. Wireless infrastructures increasingly represent strategic national assets associated with economic competitiveness, cybersecurity resilience, and technological sovereignty. Federated learning systems operating across multinational communication ecosystems may introduce concerns regarding foreign influence, infrastructure dependency, or collaborative intelligence leakage [74]. Governments may therefore impose restrictions on federated coordination across geopolitical boundaries or critical infrastructure sectors.

Spectrum regulation itself may require structural adaptation under intelligent allocation paradigms. Traditional licensing frameworks assume relatively static allocation relationships between operators and frequency bands. Adaptive reinforcement learning systems, however, continuously reconfigure spectrum usage according to real-time environmental conditions [75]. Regulators must therefore reconsider how licensing, interference management, and operational compliance are enforced within dynamically adaptive wireless ecosystems.

Standardization also becomes increasingly important. Interoperability across heterogeneous wireless infrastructures requires consistent protocols for federated synchronization, trust validation, security coordination, and policy exchange. Without standardized coordination mechanisms, fragmented proprietary implementations could undermine collaborative intelligence effectiveness and increase systemic vulnerability [76].

Fairness governance presents another significant challenge. Spectrum allocation decisions increasingly influence social inclusion, economic participation, and regional development opportunities. Federated reinforcement learning systems optimized primarily for technical efficiency may inadvertently reproduce or amplify existing infrastructural inequalities [77]. Regulatory oversight mechanisms ensuring equitable access and nondiscriminatory resource allocation therefore become essential components of responsible deployment.

Cybersecurity governance is similarly critical. Federated learning infrastructures introduce novel attack surfaces associated with distributed coordination mechanisms, model aggregation processes, and collaborative intelligence exchange. Malicious actors targeting federated wireless systems could destabilize communication infrastructures or manipulate resource allocation policies [78]. Regulatory frameworks governing cybersecurity certification, trust validation, and infrastructure resilience must therefore evolve alongside intelligent wireless technologies.

Ethical considerations further extend beyond purely technical governance. Autonomous spectrum allocation systems increasingly mediate societal access to digital services, educational platforms, healthcare systems, and economic opportunities. Decisions regarding communication prioritization, congestion management, and service differentiation consequently carry ethical implications regarding inclusion, accessibility, and distributive fairness [79].

The governance implications of federated deep reinforcement learning therefore reflect broader societal debates concerning the role of artificial intelligence in critical infrastructure management. Communication systems are no longer merely technical utilities but foundational societal coordination platforms influencing economic productivity, political stability, social interaction, and emergency resilience.

International coordination also becomes increasingly necessary because wireless ecosystems operate across interconnected global infrastructures. Harmonized standards for privacy protection, security validation, interoperability management, and ethical AI deployment may help reduce fragmentation while supporting collaborative innovation [80]. Nevertheless, achieving international consensus remains challenging amid geopolitical competition and divergent regulatory philosophies.

Ultimately, governance frameworks for intelligent spectrum allocation must balance innovation enablement against societal protection. Excessively restrictive regulation may slow technological advancement, while insufficient oversight could permit systemic vulnerabilities, discriminatory behaviors, or infrastructure instability. Federated deep reinforcement learning therefore requires governance models capable of integrating technical flexibility with institutional accountability and public-interest safeguards.

8. Deployment Scenarios and Cross-Domain Applications

The practical relevance of federated deep reinforcement learning for adaptive spectrum allocation becomes increasingly apparent when examining diverse deployment scenarios across contemporary wireless ecosystems. Dense 5G-A infrastructures support a wide range of operational environments characterized by distinct performance requirements, mobility dynamics, governance constraints, and infrastructural conditions. Consequently, adaptive spectrum allocation frameworks must demonstrate flexibility across heterogeneous communication domains rather than functioning solely within idealized simulation environments [81].

Urban smart city environments represent one of the most prominent deployment contexts. Metropolitan wireless ecosystems contain overlapping layers of public communication infrastructure, transportation systems, environmental monitoring networks, surveillance platforms, emergency response services, and commercial edge applications. Traffic patterns fluctuate continuously according to commuting behavior, public events, weather conditions, and economic activity [82]. Federated reinforcement learning enables localized spectrum adaptation within individual urban districts while maintaining collaborative coordination across citywide infrastructures.

In densely populated commercial districts, communication demand may surge rapidly during peak business hours or major public gatherings. Local reinforcement learning agents embedded within edge base stations can dynamically reallocate spectrum resources according to evolving traffic conditions while federated coordination enables neighboring districts to share congestion management strategies [83]. This collaborative adaptability improves operational resilience under highly dynamic urban conditions.

Industrial automation systems present another important application domain. Smart manufacturing facilities increasingly rely on ultra-reliable low-latency communication infrastructures supporting robotic coordination, predictive maintenance, autonomous logistics, and real-time process monitoring. Industrial wireless environments often exhibit highly specialized interference characteristics and deterministic communication requirements [84].

Federated reinforcement learning offers significant advantages within industrial contexts because factories frequently prefer localized operational autonomy and data privacy preservation. Sensitive manufacturing data, proprietary process information, and operational telemetry can remain within local facilities while collaborative intelligence sharing enables broader optimization across distributed industrial networks [85]. Adaptive spectrum allocation may also improve coexistence between industrial wireless systems and surrounding public communication infrastructures.

Vehicular communication systems constitute another highly dynamic deployment environment. Intelligent transportation infrastructures involve continuous interaction among connected vehicles, roadside edge units, traffic coordination systems, and cloud-based mobility platforms. Spectrum allocation within vehicular environments must adapt rapidly to changing mobility patterns, traffic density fluctuations, and environmental conditions [86].

Federated reinforcement learning supports mobility-aware optimization by enabling distributed agents to learn from geographically diverse traffic behaviors and road conditions. Highway communication infrastructures, urban intersections, and logistics corridors may collaboratively refine spectrum allocation policies without centralized dependence [87]. This distributed adaptability improves responsiveness under high-speed mobility scenarios where centralized coordination latency may become problematic.

Healthcare communication systems also represent increasingly important deployment contexts. Hospitals, telemedicine platforms, remote diagnostics infrastructures, and emergency response networks require highly reliable wireless communication capabilities. Adaptive spectrum allocation can prioritize latency-sensitive medical applications while mitigating congestion during large-scale healthcare emergencies or disaster response operations [88].

Privacy preservation is particularly important within healthcare environments because communication data may contain sensitive patient information or operational details regarding critical medical infrastructures. Federated learning therefore aligns well with healthcare communication governance requirements by minimizing centralized exposure of sensitive operational datasets [89].

Rural and underserved regions present distinct deployment considerations. These areas often experience infrastructure scarcity, limited computational resources, and intermittent connectivity conditions. Federated reinforcement learning architectures must therefore operate efficiently under constrained infrastructural environments rather than assuming ideal edge computing availability [90].

Adaptive spectrum allocation may nevertheless improve communication accessibility within underserved regions by dynamically optimizing scarce radio resources according to local demand conditions. Collaborative learning across geographically distributed rural infrastructures can also accelerate knowledge transfer and operational adaptation despite limited local data availability [91].

Military and defense communication systems represent another strategically significant domain. Tactical communication environments require resilient spectrum allocation under adversarial conditions, infrastructure disruption, and dynamic operational mobility. Federated reinforcement learning supports decentralized coordination capabilities that may improve communication continuity during contested or degraded operational scenarios [92].

However, defense-oriented deployments also intensify governance and security concerns because adversarial manipulation of collaborative learning processes could compromise operational effectiveness. Trust-aware coordination mechanisms and robust cybersecurity validation therefore become especially critical within such environments [93].

Large-scale public events and disaster response operations further illustrate the importance of adaptive spectrum allocation. Communication demand may increase dramatically during emergencies, infrastructure failures, or mass gatherings. Federated reinforcement learning enables distributed infrastructures to collaboratively adapt allocation strategies according to rapidly evolving situational conditions [94]. Decentralized coordination also improves resilience when centralized communication facilities become inaccessible or overloaded.

Cross-domain deployment experiences collectively demonstrate that federated deep reinforcement learning should not be viewed as a universally uniform solution. Effective implementation depends heavily on contextual adaptation according to operational environments, governance expectations, infrastructural conditions, and societal priorities. The broader significance of adaptive spectrum allocation therefore lies not only in technical optimization but also in enabling resilient, context-sensitive, and socially responsive wireless ecosystems.

9. Future Research Directions and Emerging Infrastructure Paradigms

The future evolution of federated deep reinforcement learning for adaptive spectrum allocation will likely be shaped by broader transformations in wireless infrastructure design, distributed intelligence architectures, and socio-technical governance models. As communication ecosystems continue evolving toward increasingly autonomous, immersive, and interconnected environments, intelligent spectrum management frameworks must adapt to emerging operational complexities and societal expectations [95].

One important research direction involves integrating semantic communication principles into federated wireless intelligence architectures. Conventional communication systems prioritize accurate data transmission regardless of contextual relevance, whereas semantic communication emphasizes the transmission of meaningful information according to application requirements. Adaptive spectrum allocation mechanisms may increasingly optimize communication resources based on semantic importance rather than purely quantitative traffic metrics [96].

The emergence of sixth-generation communication research further expands the conceptual scope of intelligent wireless systems. Future infrastructures are expected to support integrated sensing and communication, holographic interaction environments, distributed digital twins, autonomous robotic ecosystems, and pervasive ambient intelligence. These applications will likely generate unprecedented demands for adaptive spectrum coordination and real-time distributed intelligence [97].

Federated reinforcement learning may therefore evolve toward hierarchical multi-agent coordination frameworks capable of managing highly heterogeneous communication ecosystems across terrestrial, aerial, maritime, and satellite infrastructures. Integrated space-air-ground communication systems introduce multidimensional mobility and propagation dynamics that substantially exceed the complexity of current 5G-A environments [98].

Digital twin integration represents another promising research trajectory. Wireless infrastructure digital twins enable real-time virtual representations of communication environments capable of supporting predictive optimization, anomaly detection, and policy simulation. Federated reinforcement learning agents may increasingly interact with digital twin environments to accelerate policy training and improve adaptation under rare or extreme operational scenarios [99].

Explainable artificial intelligence will also become increasingly important. Regulatory institutions, infrastructure operators, and public stakeholders are likely to demand greater transparency regarding how adaptive spectrum allocation decisions are generated. Future federated reinforcement learning systems may therefore incorporate interpretable policy mechanisms capable of supporting operational auditing and governance accountability [100].

Trust management frameworks require further development as well. Collaborative wireless intelligence ecosystems depend heavily on reliable coordination among distributed participants. Blockchain-enabled trust validation, decentralized identity management, and reputation-aware aggregation mechanisms may help strengthen the integrity of federated learning processes under adversarial conditions [101].

Quantum communication and quantum computing research may also influence future wireless intelligence paradigms. While practical large-scale quantum wireless systems remain speculative, quantum-enhanced optimization methods could potentially improve distributed resource allocation efficiency under extremely complex operational conditions [102]. However, such developments would also introduce new cybersecurity and governance challenges.

Environmental sustainability will likely remain a defining research priority. Future communication systems must balance increasing computational intensity against global carbon reduction objectives and resource sustainability expectations. Energy-aware federated reinforcement learning, low-power edge intelligence, and carbon-adaptive communication orchestration may consequently become central research themes [103].

Human-centered infrastructure governance also warrants greater attention. Intelligent spectrum allocation systems increasingly mediate social participation, economic access, and institutional coordination. Future research should therefore examine how adaptive wireless infrastructures influence societal equity, digital inclusion, labor transformation, and democratic accountability [104].

Cross-disciplinary collaboration will become increasingly necessary for addressing these challenges effectively. Wireless communication engineering alone cannot fully resolve the governance, ethical, environmental, and societal implications associated with autonomous spectrum allocation systems. Meaningful progress will require coordinated contributions from computer science, public policy, economics, law, environmental science, and social systems research communities [105].

The long-term trajectory of federated deep reinforcement learning ultimately reflects a broader transformation toward intelligent infrastructure ecosystems capable of collaborative self-organization under conditions of uncertainty and complexity. Adaptive spectrum allocation may therefore become one component within larger autonomous infrastructure frameworks encompassing transportation, energy systems, healthcare coordination, industrial automation, and urban governance.

10. Conclusion

The increasing complexity of dense 5G-A wireless ecosystems has transformed adaptive spectrum allocation into a multidimensional infrastructure challenge involving not only communication efficiency but also resilience, fairness, sustainability, governance, and societal coordination. Traditional centralized spectrum management approaches are increasingly insufficient for addressing the dynamic operational conditions generated by edge-native intelligence, heterogeneous service demands, large-scale device connectivity, and rapidly evolving radio environments.

Federated deep reinforcement learning represents a promising paradigm for enabling distributed and collaborative spectrum intelligence within these complex communication ecosystems. By combining adaptive reinforcement learning capabilities with privacy-preserving federated coordination mechanisms, wireless infrastructures can achieve more responsive, scalable, and resilient spectrum allocation strategies while reducing dependence on centralized operational control.

This paper has examined the architectural foundations, operational trade-offs, deployment challenges, governance implications, and sustainability considerations associated with federated reinforcement learning in dense 5G-A networks. The analysis demonstrates that distributed collaborative intelligence offers substantial advantages for managing interference complexity, mobility-aware optimization, network slicing coordination, and localized adaptation under highly dynamic conditions.

At the same time, federated deep reinforcement learning introduces significant systemic challenges involving scalability, model convergence, cybersecurity resilience, fairness preservation, and environmental sustainability. Effective deployment therefore requires careful integration of trust-aware coordination frameworks, explainable governance mechanisms, fairness-sensitive optimization strategies, and energy-efficient infrastructure design principles.

The broader significance of adaptive spectrum allocation extends beyond technical optimization alone. Intelligent wireless infrastructures increasingly influence economic productivity, healthcare accessibility, transportation coordination, industrial automation, emergency response capabilities, and societal resilience. Consequently, federated spectrum intelligence must be evaluated through holistic socio-technical perspectives encompassing institutional accountability, public-interest governance, and long-term sustainability objectives.

Future wireless ecosystems will likely become even more decentralized, autonomous, and interconnected as edge intelligence, digital twins, integrated sensing systems, and sixth-generation communication paradigms continue to evolve. Federated deep reinforcement learning provides a foundational framework for supporting this transition toward adaptive infrastructure intelligence, particularly when combined with robust governance models and interdisciplinary systems-oriented research approaches.

Ultimately, the future success of intelligent spectrum allocation will depend not only on advances in machine learning architectures but also on the development of trustworthy, equitable, sustainable, and resilient communication ecosystems capable of supporting increasingly network-dependent societies.

References

[1] Andrews, J. G., Buzzi, S., Choi, W., Hanly, S. V., Lozano, A., Soong, A. C., & Zhang, J. C. (2014). What will 5G be? IEEE Journal on Selected Areas in Communications, 32(6), 1065–1082.

[2] Akyildiz, I. F., Nie, S., Lin, S. C., & Chandrasekaran, M. (2016). 5G roadmap: 10 key enabling technologies. Computer Networks, 106, 17–48.

[3] Zhang, Y., Wang, H., Zheng, T., & Yang, Q. (2018). Deep reinforcement learning for resource management in network slicing. IEEE Access, 6, 74429–74441.

[4] Mao, Y., You, C., Zhang, J., Huang, K., & Letaief, K. B. (2017). A survey on mobile edge computing: The communication perspective. IEEE Communications Surveys & Tutorials, 19(4), 2322–2358.

[5] Mnih, V., Kavukcuoglu, K., Silver, D., Rusu, A. A., Veness, J., Bellemare, M. G., & Hassabis, D. (2015). Human-level control through deep reinforcement learning. Nature, 518(7540), 529–533.

[6] Wang, S., Tuor, T., Salonidis, T., Leung, K. K., Makaya, C., He, T., & Chan, K. (2019). Adaptive federated learning in resource constrained edge computing systems. IEEE Journal on Selected Areas in Communications, 37(6), 1205–1221.

[7] McMahan, H. B., Moore, E., Ramage, D., Hampson, S., & Aguera y Arcas, B. (2017). Communication-efficient learning of deep networks from decentralized data. Proceedings of AISTATS, 1273–1282.

[8] Nguyen, D. C., Ding, M., Pathirana, P. N., Seneviratne, A., Li, J., Niyato, D., & Poor, H. V. (2020). Federated learning for wireless communications: Motivation, opportunities, and challenges. IEEE Communications Surveys & Tutorials, 22(4), 2642–2683.

[9] Ziegeldorf, J. H., Morchon, O. G., & Wehrle, K. (2014). Privacy in the Internet of Things: Threats and challenges. Security and Communication Networks, 7(12), 2728–2742.

[10] Haykin, S. (2005). Cognitive radio: Brain-empowered wireless communications. IEEE Journal on Selected Areas in Communications, 23(2), 201–220.

[11] Dahlman, E., Parkvall, S., & Skold, J. (2018). 5G NR: The next generation wireless access technology. Academic Press.

[12] Shafi, M., Molisch, A. F., Smith, P. J., Haustein, T., Zhu, P., Silva, P. D., & Wunder, G. (2017). 5G: A tutorial overview of standards, trials, challenges, deployment, and practice. IEEE Journal on Selected Areas in Communications, 35(6), 1201–1221.

[13] Boccardi, F., Heath, R. W., Lozano, A., Marzetta, T. L., & Popovski, P. (2014). Five disruptive technology directions for 5G. IEEE Communications Magazine, 52(2), 74–80.

[14] Sutton, R. S., & Barto, A. G. (2018). Reinforcement learning: An introduction. MIT Press.

[15] Li, Y. (2018). Deep reinforcement learning: An overview. arXiv preprint arXiv:1810.06339.

[16] Chen, M., Challita, U., Saad, W., Yin, C., & Debbah, M. (2017). Artificial neural networks-based machine learning for wireless networks: A tutorial. IEEE Communications Surveys & Tutorials, 21(4), 3039–3071.

[17] Lim, W. Y. B., Luong, N. C., Hoang, D. T., Jiao, Y., Liang, Y. C., Yang, Q., & Miao, C. (2020). Federated learning in mobile edge networks: A comprehensive survey. IEEE Communications Surveys & Tutorials, 22(3), 2031–2063.

[18] Ye, H., Li, G. Y., & Juang, B. H. (2019). Deep reinforcement learning based resource allocation for V2V communications. IEEE Transactions on Vehicular Technology, 68(4), 3163–3173.

[19] Park, J., Samarakoon, S., Bennis, M., & Debbah, M. (2019). Wireless network intelligence at the edge. Proceedings of the IEEE, 107(11), 2204–2239.

[20] Saad, W., Bennis, M., & Chen, M. (2020). A vision of 6G wireless systems: Applications, trends, technologies, and open research problems. IEEE Network, 34(3), 134–142.

[21] Nasir, A. A., Guo, S., & Wang, J. (2021). Multi-agent deep reinforcement learning for dynamic spectrum access in multichannel wireless networks. IEEE Transactions on Cognitive Communications and Networking, 7(1), 33–45.

[22] Kairouz, P., McMahan, H. B., Avent, B., Bellet, A., Bennis, M., Bhagoji, A. N., & Zhao, S. (2021). Advances and open problems in federated learning. Foundations and Trends in Machine Learning, 14(1–2), 1–210.

[23] Mach, P., & Becvar, Z. (2017). Mobile edge computing: A survey on architecture and computation offloading. IEEE Communications Surveys & Tutorials, 19(3), 1628–1656.

[24] Abbas, N., Zhang, Y., Taherkordi, A., & Skeie, T. (2018). Mobile edge computing: A survey. IEEE Internet of Things Journal, 5(1), 450–465.

[25] Roman, R., Lopez, J., & Mambo, M. (2018). Mobile edge computing, fog et al.: A survey and analysis of security threats and challenges. Future Generation Computer Systems, 78, 680–698.

[26] Zhao, Y., Li, M., Lai, L., Suda, N., Civin, D., & Chandra, V. (2018). Federated learning with non-IID data. arXiv preprint arXiv:1806.00582.

[27] Bonawitz, K., Eichner, H., Grieskamp, W., Huba, D., Ingerman, A., Ivanov, V., & Van Overveldt, T. (2019). Towards federated learning at scale: System design. Proceedings of MLSys, 374–388.

[28] Samarakoon, S., Bennis, M., Saad, W., & Debbah, M. (2019). Distributed federated learning for ultra-reliable low-latency vehicular communications. IEEE Transactions on Communications, 68(2), 1146–1159.

[29] Bagdasaryan, E., Veit, A., Hua, Y., Estrin, D., & Shmatikov, V. (2020). How to backdoor federated learning. Proceedings of AISTATS, 2938–2948.

[30] Foukas, X., Patounas, G., Elmokashfi, A., & Marina, M. K. (2017). Network slicing in 5G: Survey and challenges. IEEE Communications Magazine, 55(5), 94–100.

[31] Wang, X., Han, Y., Wang, V. C. M., Zhao, Q., Chen, X., & Chen, M. (2020). In-edge AI: Intelligentizing mobile edge computing, caching and communication by federated learning. IEEE Network, 33(5), 156–165.

[32] Challita, U., Ferdowsi, A., Chen, M., & Saad, W. (2019). Machine learning for wireless connectivity and security of cellular-connected UAVs. IEEE Wireless Communications, 26(1), 28–35.

[33] Luong, N. C., Hoang, D. T., Gong, S., Niyato, D., Wang, P., Liang, Y. C., & Kim, D. I. (2019). Applications of deep reinforcement learning in communications and networking. IEEE Communications Surveys & Tutorials, 21(4), 3133–3174.

[34] Zhou, Z., Chen, X., Li, E., Zeng, L., Luo, K., & Zhang, J. (2019). Edge intelligence: Paving the last mile of artificial intelligence with edge computing. Proceedings of the IEEE, 107(8), 1738–1762.

[35] Li, Q. (2026). QoS Assurance Mechanism for 5G Network Slicing Based on the Deep Reinforcement Learning PPO Algorithm. arXiv preprint arXiv:2605.03345.

[36] Sun, Y., Peng, M., Mao, S., Wang, Y., & Liu, C. (2018). Application of machine learning in wireless networks: Key techniques and open issues. IEEE Communications Surveys & Tutorials, 21(4), 3072–3108.

[37] Campolo, C., Molinaro, A., Araniti, G., & Berthet, A. O. (2017). Better platooning control toward autonomous driving: An LTE device-to-device communications strategy. Vehicular Communications, 16, 1–10.

[38] Shi, W., Cao, J., Zhang, Q., Li, Y., & Xu, L. (2016). Edge computing: Vision and challenges. IEEE Internet of Things Journal, 3(5), 637–646.

[39] Wang, C., Wang, C., Liang, Y. C., Chen, F. R., & Cheng, H. (2020). Learning-driven resource allocation for wireless communication systems. IEEE Wireless Communications, 27(4), 10–16.

[40] Rawls, J. (1971). A theory of justice. Harvard University Press.

[41] Nishio, T., & Yonetani, R. (2019). Client selection for federated learning with heterogeneous resources in mobile edge. Proceedings of ICC, 1–7.

[42] Zhu, L., Liu, Z., & Han, S. (2019). Deep leakage from gradients. Advances in Neural Information Processing Systems, 32, 14774–14784.

[43] Wang, J., Liu, Q., Liang, H., Joshi, G., & Poor, H. V. (2020). Tackling the objective inconsistency problem in heterogeneous federated optimization. Advances in Neural Information Processing Systems, 33, 7611–7623.

[44] Letaief, K. B., Shi, Y., Lu, J., & Lu, J. (2019). Edge artificial intelligence for 6G: Vision, enabling technologies, and applications. IEEE Journal on Selected Areas in Communications, 40(1), 5–36.

[45] Deng, Y., Lyu, F., Ren, J., Yang, H., Zhou, Y., Zhang, Y., & Shen, X. (2020). FAIR: Quality-aware federated learning for green mobile edge networks. Proceedings of ICC, 1–6.

[46] Hsieh, K., Phanishayee, A., Mutlu, O., & Gibbons, P. B. (2020). The non-IID data quagmire of decentralized machine learning. Proceedings of ICML, 4387–4398.

[47] Tran, N. H., Bao, W., Zomaya, A., Nguyen, M. N. H., & Hong, C. S. (2019). Federated learning over wireless networks: Optimization model design and analysis. Proceedings of INFOCOM, 1387–1395.

[48] Yu, W., Liang, F., He, X., Hatcher, W. G., Lu, C., Lin, J., & Yang, X. (2018). A survey on the edge computing for the Internet of Things. IEEE Access, 6, 6900–6919.

[49] Fung, C., Yoon, C. J. M., & Beschastnikh, I. (2020). The limitations of federated learning in sybil settings. Proceedings of RAID, 301–316.

[50] Shi, Y., Sagduyu, Y. E., Davaslioglu, K., Li, J. H., & Erpek, T. (2021). Adversarial deep learning for cognitive radio security. IEEE Wireless Communications, 28(3), 64–71.

[51] Sutton, R. S. (1990). Integrated architectures for learning, planning, and reacting based on approximating dynamic programming. Proceedings of ICML, 216–224.

[52] van Dijk, J. (2020). The digital divide. Polity Press.

[53] Eubanks, V. (2018). Automating inequality. St. Martin's Press.

[54] Helou, M. E., Ibrahim, S., & Machanavajjhala, A. (2020). The price of fairness in federated learning. Proceedings of ICML Workshops.

[55] Barocas, S., Hardt, M., & Narayanan, A. (2019). Fairness and machine learning. fairmlbook.org.

[56] Strubell, E., Ganesh, A., & McCallum, A. (2019). Energy and policy considerations for deep learning in NLP. Proceedings of ACL, 3645–3650.

[57] Bianzino, A. P., Chaudet, C., Rossi, D., & Rougier, J. L. (2012). A survey of green networking research. IEEE Communications Surveys & Tutorials, 14(1), 3–20.

[58] Jones, N. (2018). How to stop data centres from gobbling up the world's electricity. Nature, 561(7722), 163–166.

[59] Henderson, P., Hu, J., Romoff, J., Brunskill, E., Jurafsky, D., & Pineau, J. (2018). Deep reinforcement learning that matters. Proceedings of AAAI, 3207–3214.

[60] Ismail, M., Zhuang, W., Serpedin, E., & Qaraqe, K. (2015). A survey on green mobile networking: From the perspectives of network operators and mobile users. IEEE Communications Surveys & Tutorials, 17(3), 1535–1556.

[61] Li, T., Sahu, A. K., Talwalkar, A., & Smith, V. (2020). Federated learning: Challenges, methods, and future directions. IEEE Signal Processing Magazine, 37(3), 50–60.

[62] Xu, D., Li, P., Guo, S., & Wu, S. (2021). Deep reinforcement learning enabled green heterogeneous networks. IEEE Transactions on Green Communications and Networking, 5(3), 1109–1121.

[63] Han, T., & Ansari, N. (2018). Green mobile edge computing for 5G and beyond. IEEE Communications Magazine, 56(8), 123–129.

[64] Chien, S. Y., Hsiang, H. P., Hsu, C. Y., & Wang, K. C. (2020). Deep learning-based energy-efficient resource allocation for future wireless networks. IEEE Access, 8, 56813–56827.

[65] Parajuly, K., Kuehr, R., Awasthi, A. K., Fitzpatrick, C., Lepawsky, J., Smith, E., & Zeng, X. (2019). Future e-waste scenarios. United Nations University.

[66] Deng, S., Zhao, H., Fang, W., Yin, J., Dustdar, S., & Zomaya, A. Y. (2020). Edge intelligence: The confluence of edge computing and artificial intelligence. IEEE Internet of Things Journal, 7(8), 7457–7469.

[67] IPCC. (2022). Climate change 2022: Impacts, adaptation and vulnerability. Cambridge University Press.

[68] Sen, A. (2009). The idea of justice. Harvard University Press.

[69] European Commission. (2020). A new circular economy action plan. Brussels.

[70] Cave, M., & Webb, W. (2015). Spectrum management: Using the airwaves for maximum social and economic benefit. Cambridge University Press.

[71] Floridi, L., & Cowls, J. (2019). A unified framework of five principles for AI in society. Harvard Data Science Review, 1(1), 1–15.

[72] Doshi-Velez, F., & Kim, B. (2017). Towards a rigorous science of interpretable machine learning. arXiv preprint arXiv:1702.08608.

[73] Shokri, R., & Shmatikov, V. (2015). Privacy-preserving deep learning. Proceedings of CCS, 1310–1321.

[74] Nye, J. S. (2011). The future of power. Public Affairs.

[75] Weiss, M. B., Lehr, W. H., & Cui, Y. (2015). Technological change and spectrum governance. Telecommunications Policy, 39(10–11), 824–836.

[76] 3GPP. (2022). Study on AI/ML for NR air interface. Technical Specification Group Radio Access Network.

[77] Noble, S. U. (2018). Algorithms of oppression. NYU Press.

[78] ENISA. (2021). Threat landscape for 5G networks. European Union Agency for Cybersecurity.

[79] Winner, L. (1980). Do artifacts have politics? Daedalus, 109(1), 121–136.

[80] OECD. (2019). OECD principles on artificial intelligence. OECD Publishing.

[81] Cisco. (2023). Annual internet report. Cisco Systems.

[82] Batty, M. (2018). Inventing future cities. MIT Press.

[83] Zanella, A., Bui, N., Castellani, A., Vangelista, L., & Zorzi, M. (2014). Internet of Things for smart cities. IEEE Internet of Things Journal, 1(1), 22–32.

[84] Lu, Y. (2017). Industry 4.0: A survey on technologies, applications and open research issues. Journal of Industrial Information Integration, 6, 1–10.

[85] Xu, X., Lu, Y., Vogel-Heuser, B., & Wang, L. (2021). Industry 4.0 and Industry 5.0: Inception, conception and perception. Journal of Manufacturing Systems, 61, 530–535.

[86] Gerla, M., Lee, E. K., Pau, G., & Lee, U. (2014). Internet of vehicles: From intelligent grid to autonomous cars and vehicular clouds. Proceedings of WF-IoT, 241–246.

[87] Hartenstein, H., & Laberteaux, K. P. (2010). VANET: Vehicular applications and inter-networking technologies. Wiley.

[88] Islam, S. M. R., Kwak, D., Kabir, M. H., Hossain, M., & Kwak, K. S. (2015). The Internet of Things for health care: A comprehensive survey. IEEE Access, 3, 678–708.

[89] Rieke, N., Hancox, J., Li, W., Milletari, F., Roth, H. R., Albarqouni, S., & Cardoso, M. J. (2020). The future of digital health with federated learning. NPJ Digital Medicine, 3(119), 1–7.

[90] Townsend, A. M. (2013). Smart cities. W. W. Norton & Company.

[91] Graham, M. (2019). Digital economies at global margins. MIT Press.

[92] Taddeo, M., & Floridi, L. (2018). How AI can be a force for good. Science, 361(6404), 751–752.

[93] Scharre, P. (2018). Army of none: Autonomous weapons and the future of war. W. W. Norton & Company.

[94] Comfort, L. K. (2007). Crisis management in hindsight. Public Administration Review, 67, 189–197.

[95] Latva-aho, M., & Leppanen, K. (2019). Key drivers and research challenges for 6G ubiquitous wireless intelligence. University of Oulu.

[96] Strinati, E. C., Barbarossa, S., Gonzalez-Jimenez, J. L., Ktenas, D., Cassau, J., Maret, L., & Dehos, C. (2021). 6G: The next frontier. arXiv preprint arXiv:1901.03239.

[97] Dang, S., Amin, O., Shihada, B., & Alouini, M. S. (2020). What should 6G be? Nature Electronics, 3(1), 20–29.

[98] Zhang, N., Yang, P., Ren, J., Shen, X., & Mark, J. W. (2019). Software defined space-air-ground integrated vehicular networks. IEEE Communications Magazine, 57(7), 50–56.

[99] Tao, F., Zhang, H., Liu, A., & Nee, A. Y. C. (2019). Digital twin in industry: State-of-the-art. IEEE Transactions on Industrial Informatics, 15(4), 2405–2415.

[100] Miller, T. (2019). Explanation in artificial intelligence: Insights from the social sciences. Artificial Intelligence, 267, 1–38.

[101] Lu, Y., Huang, X., Dai, Y., Maharjan, S., & Zhang, Y. (2020). Blockchain and federated learning for privacy-preserved data sharing in industrial IoT. IEEE Transactions on Industrial Informatics, 16(6), 4177–4186.

[102] Gyongyosi, L., & Imre, S. (2019). A survey on quantum computing technology. Computer Science Review, 31, 51–71.

[103] Wu, Y., Khisti, A., Xiao, C., Caire, G., Wong, K. K., & Gao, X. (2018). A survey of physical layer security techniques for 5G wireless networks. IEEE Journal on Selected Areas in Communications, 36(4), 679–695.

[104] Couldry, N., & Mejias, U. A. (2019). The costs of connection. Stanford University Press.

[105] Jasanoff, S. (2016). The ethics of invention. W. W. Norton & Company.