

AI Assisted Resource Management for Private 5G NetworksRifki Alfairizi^{*}, Aditya²,*1 Teknik Elektro, Fakultas Teknik, Garut University, Indonesia**2 Teknik Elektro, Fakultas Teknik, Garut University, Indonesia*

alfairizi2009@gmail.com, aditya_yayang@gmail.com

Received	10 01 2026	Revised	25 01 2026
Accepted	30 01 2026	Published	02 02 2026

Abstract Private fifth generation mobile networks offer industrial sites direct control over coverage, quality of service, security boundaries, and integration with operational applications. Their flexibility also creates a difficult resource management problem because radio spectrum, transport, computing, and network functions must serve traffic with different latency, reliability, throughput, and device density requirements. This technical review examines how artificial intelligence can assist resource management without weakening service guarantees. It relates 5G system architecture, network slicing, multi access edge computing, and open radio access network control to supervised learning, anomaly detection, contextual bandits, and deep reinforcement learning. A practical decision framework separates planning, prediction, optimization, admission control, and closed loop adaptation. It recommends that learning models estimate demand and radio conditions, while policy engines and constrained optimizers retain explicit responsibility for service limits. For higher autonomy, learned actions should pass through feasibility and safety checks, operate within bounded scopes, and fall back to deterministic scheduling. Unlike prior surveys that catalogue machine-learning techniques across wireless networking broadly, this review is scoped specifically to the industrial private-5G setting and to the governance boundary between learned recommendations and the deterministic policy engine that retains authority over service limits. The article identifies data requirements, digital twin and testbed needs, evaluation metrics, security controls, and a staged pathway suitable for manufacturing campuses. The main conclusion is that artificial intelligence should be evaluated against strong engineering baselines and judged by tail latency, reliability violations, stability, and operational burden rather than average throughput alone. This approach enables private 5G systems to gain adaptive efficiency while preserving auditable control.

Keywords *private 5G; radio resource management; network slicing; machine learning; O RAN; industrial communications*

1 Introduction

Private 5G networks are deployed for organizations that need local control over connectivity, data paths, and service behavior. A manufacturing campus may connect mobile robots, machine vision cameras, programmable controllers, handheld terminals, environmental sensors, and enterprise devices on the same physical infrastructure. These applications differ sharply in packet size, periodicity, mobility, latency, reliability, and consequence of failure. The 5G system provides flexible quality of service and network slicing mechanisms, but those mechanisms still require resource allocation policies that translate business and operational needs into radio, transport, core, and computing decisions [1]-[3].

Traditional radio resource management relies on scheduling rules, link adaptation, power control, admission control, and engineered margins. These methods remain essential because they have known timing behavior and can be tested against standards. Complexity grows when the network must coordinate several cells, slices, edge applications, and changing production schedules. Interference, mobility, traffic bursts, and compute contention interact. A decision that raises radio throughput can increase queueing at an edge service, while aggressive admission can reduce the

reserve required by a safety related flow. Cross domain optimization is therefore attractive but also difficult to formulate and solve in real time.

Machine learning can estimate demand, classify traffic, predict channel quality, detect anomalies, or approximate policies that would otherwise require repeated optimization. Research surveys describe extensive use of learning in wireless networking and network slicing [4]-[7]. These methods can exploit patterns that fixed rules miss, but they do not create service guarantees by themselves. Training data may not include rare faults or production transitions, and a model can change behavior when devices, software, or spectrum conditions change. Closed loop use consequently requires boundaries, observability, and a tested fallback.

This review develops a deployment oriented framework for AI assisted resource management in private 5G. It focuses on industrial sites with local spectrum or operator provided private service, on premises or nearby edge computing, and mixed quality of service. The article does not present synthetic performance as field evidence. Instead, it defines functional roles, architecture, algorithm selection, data and evaluation requirements, security constraints, and maturity stages that can guide a reproducible testbed or pilot. The aim is to make the use of AI technically accountable from planning through operation.

Existing surveys on machine learning for 5G [4]-[7] and on network slicing [8]-[11] establish the algorithmic landscape but are largely infrastructure agnostic, and O-RAN oriented reviews [12], [13], [19] concentrate on interface and controller design rather than on how an industrial operator should bound learned authority over a live production network. This review does not propose a new learning algorithm. Its contribution is the deployment framework itself: an explicit separation of observation, prediction, decision, enforcement, and assurance roles, a policy engine positioned to reject infeasible learned actions before they reach the radio or core, and a staged evidence pathway that ties increasing AI authority to accumulated operational evidence specifically for industrial private-5G settings such as a manufacturing campus.

2 Review Methodology

This review draws on 3GPP and ITU technical specifications, IEEE and ACM survey literature on machine learning for wireless networks, network slicing, and open radio access networks, and studies of reinforcement learning and contextual bandits for radio resource management, retrieved from IEEE Xplore, ACM Digital Library, arXiv, and relevant standards bodies using combinations of the terms *private 5G*, *network slicing*, *O-RAN*, *radio resource management*, *reinforcement learning*, and *multi-access edge computing*. Standards documents were included as authoritative sources for architecture and terminology regardless of publication year; peer-reviewed technical and survey literature was prioritized from the last five to eight years to reflect the current state of AI-assisted radio resource management. Sources were included when they described a reproducible method, a documented architecture, or an evaluated failure mode relevant to resource management; sources without methodological detail, such as vendor marketing material, were excluded. Because the included sources differ in network layer, algorithm, and evaluation metric, the evidence was synthesized narratively rather than pooled statistically. Where sources disagreed about how much authority a learned component should hold, the more conservative position was adopted and converted into a design requirement in Sections 3 and 6 rather than resolved by selective citation.

3 Private 5G Resource Domains

Resource management begins with an end to end service model. At the radio layer, the network assigns physical resource blocks, modulation and coding, transmit power, beams, and scheduling opportunities. The transport layer provides bandwidth, queues, routing, and redundancy between radio units, distributed or centralized functions, the core, and applications. The core manages sessions, policy, quality of service flows, and user plane placement. The compute layer supplies processor, memory, accelerator, and storage capacity to virtualized network functions and edge applications. Treating one domain in isolation can move congestion rather than resolve it.

Table 1 End to end resource domains

Domain	Representative resources	Service risk if unmanaged
Radio	Spectrum blocks power beams scheduler time and cell capacity	Interference congestion deadline loss and unstable mobility
Transport	Bandwidth queues routes redundancy and synchronization	Queue growth jitter packet loss and weak resilience
Core	Sessions policies quality of service flows and user plane placement	Admission failure wrong treatment and long data paths
Edge compute	Processor memory accelerators storage and application queues	Inference delay dropped work and service saturation
Operations	Configuration change windows staff attention and recovery tools	Uncontrolled change and slow incident response

Network slicing provides a logical structure for separating services over shared infrastructure. A slice can define performance, isolation, management, and lifecycle expectations, while 5G quality of service identifiers and flows support packet treatment [1], [2]. Slicing does not imply unlimited resources or perfect isolation. The operator must decide admission, reservation, sharing, scaling, and priority when demands conflict. Research on 5G slicing highlights software defined networking and network function virtualization as key enablers, while also identifying orchestration and assurance as central challenges [8]-[11].

Industrial traffic should be classified by consequence and temporal behavior rather than by application label alone. A video inspection stream can tolerate brief quality reduction but consumes high bandwidth. A mobile robot command stream may use less bandwidth yet require bounded latency and rapid recovery, reflecting the design principles of ultra reliable low latency communication [17]. Periodic sensors may tolerate delay but arrive in large numbers. Maintenance downloads can be deferred. Each class needs a service level objective, measurement point, time window, and failure response. Without this contract, an optimizer cannot distinguish a harmless throughput change from a production risk.

Private deployment models also affect available control. A fully isolated network offers direct ownership of spectrum access, core functions, and data. A shared radio or operator integrated model can reduce cost but limits the enterprise control surface. The AI function should be placed only where the organization can obtain reliable telemetry and enforce approved actions. Interfaces, update periods, and authority vary across vendors. A portable design therefore defines intent and measurements independently of a specific controller and then maps them to the available management and radio interfaces.

4 AI Assisted Control Architecture

The proposed architecture separates observation, prediction, decision, enforcement, and assurance. The observation plane collects radio counters, device context, flow statistics, edge compute load, application indicators, and configuration state. The prediction plane estimates short term demand, mobility, radio quality, and congestion risk. The decision plane evaluates candidate allocation or scaling actions. The enforcement plane converts approved decisions into vendor or standards based commands. The assurance plane verifies service outcomes, detects drift, and triggers rollback. Edge placement follows the communication and computing considerations summarized in mobile edge computing research [16], while the machine learning functions can be mapped to the lifecycle roles defined by ITU-T Recommendation Y.3172 [18]. Separating these functions makes responsibility visible and allows each component to be replaced or validated independently.

Open radio access network concepts introduce a non real time controller for policy and model lifecycle and a near real time controller for applications that influence radio behavior. This structure is attractive for learning assisted control because training, long horizon optimization, and policy definition can remain outside tight timing loops, while approved models provide inference closer to the radio network [12], [13]. A recent comprehensive survey of O-RAN deployment options reinforces that slicing-aware architectures still require an explicit arbitration layer when multiple learned applications act on overlapping radio resources [19]. A private deployment should still treat

each exposed interface as a controlled boundary. Conflicting applications, delayed telemetry, and incompatible objectives can destabilize operation if arbitration is weak.

A policy engine should define the feasible action space before a model acts. It can enforce minimum reserved capacity, admission limits, maximum handover parameter changes, edge compute headroom, and change rate constraints. The model proposes a ranked action or allocation, and a deterministic filter verifies that it respects current topology and policy. If the proposal is infeasible, stale, or out of distribution, the system uses a known scheduler or optimization baseline. This pattern converts AI from an unrestricted controller into a bounded decision aid and produces an audit trail for each intervention.

Control cadence should match the phenomenon. Capacity planning and slice dimensioning may update daily or with production plans. Traffic prediction and virtual function scaling can operate over minutes. Mobility and scheduling decisions may require seconds, sub seconds, or tighter vendor internal loops. External AI should not be inserted into a loop whose deadline, jitter, and failure behavior it cannot meet. The architecture can use slower learning functions to tune policies or reserve levels while established radio algorithms continue fast scheduling.

5 Learning and Optimization Methods

Supervised learning supports demand, mobility, and channel prediction. Useful models range from regularized regression and tree ensembles to recurrent or attention based networks. A simple model with calibrated intervals may be more valuable than a complex point predictor because the resource manager needs uncertainty to set reserve. Features can include time, production shift, device count, cell load, signal quality, handover history, and application queue measurements. Chronological validation and per service error analysis are necessary because aggregate accuracy can hide poor performance during rare but important operating states.

Unsupervised and semi supervised methods can identify anomalous traffic, radio counters, or application behavior. They are useful when labeled failures are scarce, but their outputs should not be equated directly with attacks or equipment faults. An anomaly score is evidence that the current pattern differs from the reference distribution. The assurance process should correlate it with topology, configuration changes, maintenance windows, and application events. Response can progress from increased logging to rate limiting or isolation according to confidence and consequence.

Constrained optimization remains a strong method for allocating shared resources. Mixed integer and nonlinear formulations can represent admission, placement, routing, and capacity, but large problems may not finish within an operational deadline. Decomposition, heuristics, and learned surrogates can reduce computation. A supervised model can approximate an expensive optimizer and then pass its output through a feasibility repair step. This retains a known objective while reducing online cost. The approximation error must be evaluated specifically near capacity limits, where small errors have the greatest operational effect. Federated variants of this approach have also been proposed to keep training data local to the edge while still coordinating quality-of-service decisions across cells [21].

Reinforcement learning is attractive when actions affect future load, mobility, or queue states. Deep reinforcement learning has been studied for spectrum access and resource allocation [14], [15], and subsequent survey work specifically on slicing has organized these approaches by state representation, reward design, and training regime [20]. Its strengths are sequential coordination and the ability to learn without labeled optimal actions. Its weaknesses include sample demand, reward sensitivity, unsafe exploration, and uncertain generalization. For private 5G, training should occur in simulation or from logged interaction, not through unconstrained production exploration. The reward should include explicit penalties for service violations, oscillatory changes, energy use, and rejected actions, with hard constraints enforced outside the reward.

Contextual bandits can provide a lower risk middle ground for decisions whose consequences are short lived, such as selecting among approved scheduler profiles. They learn which option works best for a context without modeling a long state trajectory. Bayesian optimization can tune a small set of parameters during planned experiments. Both methods operate over a restricted action catalog

and can be easier to govern than a fully learned controller. Method selection should reflect the time horizon, safety consequence, observability, and number of decisions rather than the perceived sophistication of the algorithm.

Table 2 Method selection for AI assisted decisions

Method	Best suited task	Assurance requirement
Supervised prediction	Demand mobility radio quality and compute load forecasting	Chronological validation calibrated uncertainty and drift checks
Anomaly detection	Telemetry traffic and configuration deviation screening	Context correlation and graduated response
Constrained optimization	Admission placement routing reservation and scaling	Feasibility deadline and solver fallback
Contextual bandit	Selection among approved short horizon policies	Restricted catalog conservative exploration and rollback
Reinforcement learning	Sequential cross domain decisions with delayed effects	Offline training hard action filter and stress testing

6 Data Engineering and Digital Evaluation

Network data must preserve temporal and semantic alignment. Counters from radio, core, transport, and applications often use different collection periods and clocks. A training record should identify network topology, software version, configuration, slice, device class, and production context. Missing samples, counter resets, delayed exports, and aggregation changes need explicit quality flags. Labels derived from tickets or alarms should record uncertainty because operational records may describe symptoms rather than root causes. Data minimization is also important; device or user identifiers should be pseudonymized when individual identity is unnecessary.

A digital evaluation environment can combine a radio simulator, traffic generator, network function emulation, and application models. It should reproduce the decisions to be tested, not attempt to mirror every physical detail. Calibration uses field measurements for path loss, traffic distributions, mobility, queue behavior, and compute service times. Validation compares simulator outputs with periods that were not used for calibration. The difference between simulation and field behavior should be reported, since a policy can exploit unrealistic assumptions in the same way it exploits an incomplete reward.

Scenario design should cover normal shifts, startup and shutdown, dense device arrival, radio interference, backhaul degradation, edge compute contention, handover concentration, software restart, and loss of telemetry. Rare service classes require targeted tests rather than waiting for natural observations. A policy should also be evaluated under traffic mixes and topology changes outside training. When the model encounters an unrecognized state, the expected outcome is conservative operation and a visible alert, not a forced prediction.

Metrics need end to end definitions. Average throughput and mean latency are insufficient for industrial assurance. Evaluation should include tail latency percentiles, packet delivery within deadline, service level violation duration, admission rejection, handover interruption, fairness, slice isolation, compute saturation, control oscillation, energy, and recovery time. Each metric must name the measurement points and aggregation window. A model may improve average utilization while increasing bursts of deadline failure, so tradeoffs should be displayed rather than compressed into a single reward value.

7 Security and Operational Assurance

AI expands the management attack surface through data pipelines, model registries, orchestration interfaces, and automated action. Telemetry integrity matters because manipulated load or radio measurements can cause harmful allocation even when the model software is unchanged. Controls include authenticated data sources, protected time synchronization, integrity checks, role separated model approval, signed deployment packages, least privilege interfaces, and immutable action logs. The architecture should prevent a training service from acquiring direct production authority.

Model specific threats include poisoning, evasion, model substitution, and extraction. Private networks may have fewer external users than public networks, but supply chain and compromised device risks remain. Defenses should prioritize architectural controls before sophisticated detection.

Training data should be versioned and traceable, model artifacts should be signed, and deployment should verify both model and feature schema. Input range and drift checks can block clearly abnormal conditions. Sensitive models or topology data should be encrypted at rest and transferred through controlled channels.

Operational assurance also requires conflict management among applications. Two control applications can recommend incompatible actions, such as increasing radio resources for one slice while a power management application seeks to reduce consumption. A central arbitration policy should establish priorities, budgets, and tie breaking rules. Changes should be rate limited and reversible. Before activation, each application should declare required inputs, affected parameters, expected cadence, constraint set, and fallback. These declarations support testing and make controller interactions reviewable.

Human operators need decision relevant explanations. For prediction, this can include recent error, uncertainty, and the main context changes. For allocation, the interface should show which service objectives are binding, what resources moved, and why an alternative was rejected. Explanations do not have to reveal internal model weights. They must help an engineer decide whether the action is plausible and safe. Alarm design should separate data failure, model drift, optimization infeasibility, enforcement failure, and service degradation because each requires a different response.

8 Deployment Roadmap

The first stage documents service classes, resource ownership, interfaces, and deterministic baselines. The network should already expose reliable measurements and maintain configuration history. The second stage applies learning to prediction and anomaly detection without changing allocation. This builds evidence about data quality and drift. The third stage runs candidate decisions in shadow mode and compares them with the existing scheduler or orchestrator. The fourth stage allows bounded changes for a noncritical slice during selected periods. Later stages extend scope only when rollback and service assurance remain effective.

Table 3 Minimum evidence for deployment stages

Stage	Minimum evidence	Control scope
Service baseline	Defined objectives measurement points and deterministic performance	Existing control
Prediction service	Stable data pipeline baseline comparison and uncertainty calibration	No allocation authority
Shadow decision	Live feasible recommendations and operator review	No enforcement
Bounded pilot	Stress tests security review rollback and noncritical slice selection	Limited parameters
Operational use	Sustained assurance change governance and periodic revalidation	Approved service domains

Acceptance criteria should be fixed before each trial. A candidate must meet all hard service limits and improve a stated primary outcome such as reserved capacity, energy, or admission efficiency. Tail behavior should remain within an agreed margin. The trial should include at least one rollback exercise and one loss of telemetry scenario. If a model improves laboratory performance but requires frequent manual intervention, its operational cost must be included. Automation is valuable only when it reduces total decision burden without obscuring risk.

Vendor portability can be improved by separating service intent from device specific commands. The organization defines performance objectives, priority, isolation, and measurement semantics in an internal model. Adapters translate that model to the deployed 5G core, radio, transport, and edge platform. This approach does not eliminate integration work, but it reduces dependence of the analytical layer on proprietary names and counters. Test cases can be reused when infrastructure changes, and discrepancies between vendor implementations become visible.

A field research program should publish topology abstractions, traffic generation methods, control intervals, baselines, software versions, and failure scenarios. Results should distinguish offline prediction, digital evaluation, testbed operation, and live production evidence. Cross site generalization should not be claimed from one campus. Instead, researchers can identify which features and constraints are site specific, then evaluate transfer with recalibration or fine tuning. These reporting practices will make AI assisted private 5G results more comparable and useful.

9 Limitations

This review is narrative and design-oriented rather than a systematic review with a registered protocol; it does not claim exhaustive retrieval of the private-5G and O-RAN literature or a formal meta-analysis, and the heterogeneity of network layers, algorithms, and evaluation metrics across the cited sources limits direct quantitative comparison. The proposed architecture and staged deployment pathway are conceptual and have not been validated on an instrumented private-5G testbed; quantitative results attributed to individual cited studies describe those studies' own reported outcomes and should not be read as evidence already obtained for the framework proposed here. Because the framework is intentionally agnostic to specific vendor interfaces, choices such as O-RAN controller vendor, core implementation, and edge platform are left to the deploying team, and these choices can materially affect achievable latency, cost, and integration effort in ways this review does not quantify.

10 Conclusion

Artificial intelligence can improve private 5G resource management by predicting demand, identifying abnormal conditions, approximating complex optimization, and selecting among approved policies. The greatest benefit appears when radio, transport, core, and edge compute resources are considered as an end to end service system. The greatest risk appears when a learned model receives broad authority without explicit service contracts, constraints, and fallback. The proposed architecture separates observation, prediction, decision, enforcement, and assurance, and places deterministic policy checks between models and network commands. Deployment should proceed from data readiness to prediction, shadow decisions, and bounded control. Success must be judged by reliability, tail latency, stability, security, and maintainability as well as utilization. This creates a credible path from research prototypes to auditable industrial operation.

AUTHOR CONTRIBUTIONS

CRediT: conceptualisation, literature synthesis, framework development, writing, review, and approval of the manuscript.

FUNDING

The authors declare that this research received no funding from any source.

CONFLICTS OF INTEREST

The authors declare that they have no conflicts of interest.

DATA AVAILABILITY STATEMENT

No original datasets were generated or analyzed in this study, as this article presents a conceptual technical review and design framework.

References

- [1] 3GPP, System architecture for the 5G System, Technical Specification 23.501, Release 17, 2024.
- [2] 3GPP, Management and orchestration concepts use cases and requirements, Technical Specification 28.530, Release 17, 2023.
- [3] ITU-R, Detailed specifications of the terrestrial radio interfaces of International Mobile Telecommunications 2020, Recommendation M.2150-1, 2022.
- [4] M. E. Morocho-Cayamcela, H. Lee, and W. Lim, Machine learning for 5G and beyond a survey, *IEEE Access*, vol. 7, pp. 137184-137206, 2019. <https://doi.org/10.1109/ACCESS.2019.2942534>
- [5] M. Chen, U. Challita, W. Saad, C. Yin, and M. Debbah, Machine learning for wireless networks with artificial intelligence a tutorial on neural networks, arXiv:1710.02913, 2017. <https://doi.org/10.48550/arXiv.1710.02913>

- [6] Y. Sun, M. Peng, Y. Zhou, Y. Huang, and S. Mao, Application of machine learning in wireless networks key techniques and open issues, *IEEE Communications Surveys and Tutorials*, vol. 21, no. 4, pp. 3072-3108, 2019. <https://doi.org/10.1109/COMST.2019.2924243>
- [7] C. Zhang, P. Patras, and H. Haddadi, Deep learning in mobile and wireless networking a survey, *IEEE Communications Surveys and Tutorials*, vol. 21, no. 3, pp. 2224-2287, 2019. <https://doi.org/10.1109/COMST.2019.2904897>
- [8] X. Foukas, G. Patounas, A. Elmokashfi, and M. K. Marina, Network slicing in 5G survey and challenges, *IEEE Communications Magazine*, vol. 55, no. 5, pp. 94-100, 2017. <https://doi.org/10.1109/MCOM.2017.1600951>
- [9] J. Ordóñez-Lucena et al., Network slicing for 5G with SDN NFV concepts architectures and challenges, *IEEE Communications Magazine*, vol. 55, no. 5, pp. 80-87, 2017. <https://doi.org/10.1109/MCOM.2017.1600935>
- [10] P. Rost et al., Network slicing to enable scalability and flexibility in 5G mobile networks, *IEEE Communications Magazine*, vol. 55, no. 5, pp. 72-79, 2017. <https://doi.org/10.1109/MCOM.2017.1600920>
- [11] I. Afolabi, T. Taleb, K. Samdanis, A. Ksentini, and H. Flinck, Network slicing and softwarization a survey on principles enabling technologies and solutions, *IEEE Communications Surveys and Tutorials*, vol. 20, no. 3, pp. 2429-2453, 2018. <https://doi.org/10.1109/COMST.2018.2815638>
- [12] M. Polese et al., Understanding O RAN architecture interfaces algorithms security and research challenges, *IEEE Communications Surveys and Tutorials*, vol. 25, no. 2, pp. 1376-1411, 2023. <https://doi.org/10.1109/COMST.2023.3239220>
- [13] L. Bonati et al., Intelligence and learning in O RAN for data driven next generation cellular networks, *IEEE Communications Magazine*, vol. 59, no. 10, pp. 21-27, 2021. <https://doi.org/10.1109/MCOM.101.2001120>
- [14] O. Naparstek and K. Cohen, Deep multi user reinforcement learning for distributed dynamic spectrum access, *IEEE Transactions on Wireless Communications*, vol. 18, no. 1, pp. 310-323, 2019. <https://doi.org/10.1109/TWC.2018.2879433>
- [15] U. Challita, L. Dong, and W. Saad, Proactive resource management for LTE in unlicensed spectrum a deep learning perspective, *IEEE Transactions on Wireless Communications*, vol. 17, no. 7, pp. 4674-4689, 2018. <https://doi.org/10.1109/TWC.2018.2830743>
- [16] Y. Mao, C. You, J. Zhang, K. Huang, and K. B. Letaief, A survey on mobile edge computing the communication perspective, *IEEE Communications Surveys and Tutorials*, vol. 19, no. 4, pp. 2322-2358, 2017. <https://doi.org/10.1109/COMST.2017.2745201>
- [17] P. Popovski et al., Wireless access for ultra reliable low latency communication principles and building blocks, *IEEE Network*, vol. 32, no. 2, pp. 16-23, 2018. <https://doi.org/10.1109/MNET.2018.1700258>
- [18] ITU-T, Architectural framework for machine learning in future networks including IMT 2020, Recommendation Y.3172, 2019.
- [19] K. Alam, M. A. Habibi, M. Tammen, D. Krummacker, W. Saad, M. Di Renzo, and H. D. Schotten, A comprehensive overview and survey of O-RAN: Exploring slicing-aware architecture, deployment options, and use cases, arXiv:2405.03555, 2024.
- [20] J. A. Hurtado Sánchez, K. Casilimas, and O. M. Caicedo Rendon, Deep reinforcement learning for resource management on network slicing: A survey, *Sensors*, vol. 22, no. 8, article 3031, 2022. <https://doi.org/10.3390/s22083031>
- [21] N. Kumar and A. Ahmad, Quality of service-aware adaptive radio resource management based on deep federated Q-learning for multi-access edge computing in beyond 5G cloud-radio access network, *Transactions on Emerging Telecommunications Technologies*, vol. 34, no. 6, article e4762, 2023. <https://doi.org/10.1002/ett.4762>