Federated Reinforcement Learning for Privacy-Preserving Resource Management in Multi-Domain 5G Network Slicing

Parth Venkataraman
Department of Computer Science, University of Alabama at Birmingham, Birmingham, AL, USA.
parth1975@uab.edu

Fangshan Hao
Department of Computer Science, George Mason University, Fairfax, VA, USA.
fangshanh@gmu.edu

Jeck J. Snath
Department of Computer Science, University of Central Florida, Orlando, FL, USA.
smithjack@ucf.edu

Abstract

The proliferation of 5G network slicing has enabled operators to instantiate multiple logical networks over a shared physical infrastructure, each tailored to distinct service requirements. However, orchestration of end-to-end slices across multiple administrative domains introduces formidable challenges in resource allocation, quality-of-service assurance, and data governance. Centralized machine learning approaches that gather sensitive traffic traces, user mobility patterns, and slice state information from participating domains raise significant privacy and regulatory concerns, particularly when data traverse jurisdictions with differing legal frameworks. Federated reinforcement learning offers a promising alternative by allowing domains to collaboratively train resource management policies without sharing raw operational data. This paper provides a comprehensive system-level analysis of federated reinforcement learning architectures designed for privacy-preserving multi-domain 5G network slicing. We examine the structural trade-offs between fully distributed and hierarchical aggregation models, the interplay between local policy generalization and global coordination, and the privacy-preserving mechanisms that underpin secure federated optimization. The discussion extends to robustness against non-identically distributed data, fairness in resource allocation across heterogeneous slice tenants, and the governance frameworks required to sustain trust among competing stakeholders. Deployment considerations encompassing edge-cloud continuum orchestration, energy efficiency, and interoperability with existing 3GPP management interfaces are critically assessed. Policy implications regarding cross-border data flow regulations, algorithmic accountability, and the evolution of service level agreements in a federated multi-domain ecosystem are explored. By synthesizing advances in federated learning, deep reinforcement learning, and 5G slicing standards, the paper articulates a research roadmap for resilient, equitable, and privacy-compliant resource management in next-generation networks.

Keywords

Federated reinforcement learning; 5G network slicing; privacy-preserving; multi-domain resource management; edge-cloud coordination; quality of service; policy optimization.

1. Introduction

Fifth-generation mobile networks have transformed telecommunications by introducing network slicing, a paradigm that enables the creation of multiple virtualized and independently managed logical networks on a common physical substrate. Each network slice can be configured to satisfy the stringent latency, bandwidth, reliability, and security requirements of diverse verticals such as autonomous driving, industrial automation, immersive media, and massive Internet of Things deployments. Achieving end-to-end slice assurance frequently requires coordination across radio access, transport, and core network segments that are operated by different infrastructure providers or administrative entities. Multi-domain slicing unlocks unprecedented flexibility but simultaneously complicates resource management, as domains must dynamically allocate spectrum, computing, and forwarding resources while respecting isolation guarantees and service level agreements that span organizational boundaries.

Traditional resource orchestration in 5G employs either static reservation policies or centralized optimization engines that collect comprehensive network telemetry from all involved domains. Centralized models, while conceptually straightforward, encounter substantial scalability bottlenecks and create single points of failure. More critically, they demand that sensitive operational data, including real-time user location, traffic patterns, and infrastructure utilization, be concentrated in a central controller. Such data aggregation conflicts with data sovereignty regulations enacted in many regions, including the General Data Protection Regulation in Europe and emerging data localization laws in Asia and the Americas. Infrastructure providers are increasingly reluctant to expose proprietary information that could enable competitors to infer their capacity plans, customer bases, or security postures. Consequently, privacy preservation has emerged as a first-order design constraint rather than an optional enhancement in multi-domain resource management.

Reinforcement learning, and deep reinforcement learning in particular, has shown remarkable effectiveness in autonomous slice lifecycle management by learning policies that map high-dimensional network states to resource configuration decisions. Nevertheless, training a high-performance centralized agent presupposes the very data pooling that multi-domain privacy constraints forbid. Federated learning offers an architectural antidote by distributing model training across participating nodes while keeping training data local. When federated optimization is combined with reinforcement learning, each domain can train its own policy or contribute to a shared policy using locally observed trajectories, sharing only model updates such as gradient vectors or parameter deltas with a coordinating server. The resulting federated reinforcement learning paradigm provides a pathway to privacy-preserving, collaborative resource management that respects data sovereignty while still harnessing cross-domain knowledge.

This paper presents an interdisciplinary examination of federated reinforcement learning as applied to multi-domain 5G network slicing. The analysis concentrates on system-level dimensions including architectural design, privacy mechanisms, structural trade-offs, governance, fairness, robustness, deployment realities, and policy implications. The aim is not to advocate for a single algorithmic recipe but to map the complex design space and to articulate the socio-technical challenges that researchers, network operators, and regulators

must collectively navigate to realize the vision of trustworthy, intelligent resource management across administrative boundaries.

2. Background and Related Work

Network slicing has been standardized by the 3rd Generation Partnership Project and extensively studied in the research community. Foundational surveys have delineated the architectural principles, enabling technologies such as software-defined networking and network function virtualization, and the orchestration challenges inherent in end-to-end slicing [3, 4, 5]. Multi-domain orchestration extends these concepts to inter-provider scenarios, requiring interfaces for federated resource negotiation and cross-domain slice stitching. Reference architectures have been proposed that introduce hierarchical orchestrators, intent-based interfaces, and blockchain-based trust mechanisms to coordinate resources across administrative boundaries [6]. However, these frameworks generally assume that domains are willing to share resource abstractions and performance monitoring data, an assumption that is increasingly untenable in competitive and regulated markets.

Concurrently, deep reinforcement learning has been widely investigated for single-domain slice resource management, with agents learning to perform admission control, spectrum allocation, virtual network function placement, and traffic steering [7]. Proximal policy optimization, actor-critic methods, and deep Q-networks have been adapted to 5G slicing use cases, demonstrating superior adaptation to traffic dynamics compared with heuristic or optimization-based approaches. When these learning methods are deployed in isolation within each domain, however, they neglect the interdependencies that arise when slices traverse multiple networks, leading to suboptimal end-to-end performance and potential violation of contractual quality-of-service guarantees.

Federated learning originated as a mechanism to enable collaborative model training on decentralized data without compromising individual data privacy [1, 2]. Early instantiations focused on supervised learning tasks for mobile keyboard prediction and healthcare, employing federated averaging to iteratively refine a global model while keeping raw data on devices. Subsequent research extended federated learning to reinforcement learning, proposing algorithms in which agents interact with their local environments and periodically synchronize policy parameters or value function approximations [8]. The federated reinforcement learning framework aligns naturally with wireless network management, where each base station or domain controller can act as a learning agent with access to local observations but benefits from the collective experience of peers facing similar traffic and radio conditions. A comprehensive survey of federated learning architectures for 5G and beyond highlights the potential for reducing communication overhead, preserving user privacy, and enabling multi-stakeholder collaboration [11]. Integration of federated reinforcement learning with deep reinforcement learning has been examined for beyond-5G systems, underscoring open problems in data heterogeneity, incentive design, and security [12].

Privacy preservation in federated learning has been addressed through techniques such as differential privacy, secure multi-party computation, and homomorphic encryption, each introducing a distinct balance between privacy guarantees, computational overhead, and model utility [9, 10]. In the context of 5G slicing, where the adversarial model may include honest-but-curious infrastructure providers as well as external attackers, the choice of privacy mechanism must consider the sensitivity of traffic matrices, slice performance counters, and user equipment traces. Recent studies have started to integrate privacy-preserving

mechanisms into wireless federated learning, although the specific requirements of multi-domain slicing orchestration remain underexplored.

3. System Architecture for Federated Reinforcement Learning in Multi-Domain Slicing

Designing a federated reinforcement learning system for multi-domain network slicing requires reconciling two distinct hierarchies: the management hierarchy of 5G orchestration and the aggregation topology of federated learning. A typical multi-domain slicing architecture comprises a set of administrative domains, each operating its own network slice subnet management function and domain orchestrator, and a cross-domain orchestrator responsible for end-to-end slice lifecycle coordination. In a fully decentralized federated architecture, the cross-domain orchestrator acts as a parameter server that aggregates policy updates received from domain-level agents without accessing raw telemetry. Domain agents interact with their local slice management systems, collecting state information such as current resource utilization, queue lengths, admission counts, and slice-specific performance metrics, and execute reinforcement learning steps using a locally defined reward function that balances domain-specific objectives with global slice commitments. After a configurable number of local training episodes, agents transmit encrypted policy gradients or model weights to the aggregation server, which computes a fused global policy and redistributes it to all participants.

A more hierarchical architecture introduces intermediate aggregation points aligned with technological or geographical groupings, such as a regional aggregation server that synthesizes updates from multiple radio access network domains before forwarding aggregated parameters to a core network domain controller. This topology reduces communication costs and mitigates the effects of statistical heterogeneity by allowing subsets of domains with similar traffic profiles to form specialized sub-models within the broader federation. However, hierarchical aggregation amplifies concerns about trust concentration, because intermediate aggregators could potentially reconstruct finer-grained information about constituent domains. The selection between flat and hierarchical topologies is therefore not merely an engineering trade-off but a governance decision that reflects the degree of mutual trust among infrastructure providers and the prevailing regulatory environment.

The interface between the reinforcement learning agents and the network infrastructure must support low-latency state retrieval and action execution without disrupting production traffic. Domain agents require access to standardized telemetry streams conforming to 3GPP management service specifications, as well as the ability to enact configuration changes through intent-based northbound interfaces. This necessitates tight integration with service-based management architectures and event exposure frameworks, ensuring that learning-driven decisions can be seamlessly translated into slice reconfiguration commands. The system must also accommodate asynchronous participation, where domains may join or leave the federation dynamically due to policy changes, network upgrades, or commercial renegotiations. The architectural resilience to such churn is essential for sustained operation in real-world multi-domain ecosystems.

4. Privacy-Preserving Mechanisms and Federated Learning Design

A core motivation for adopting federated reinforcement learning is the preservation of data confidentiality across administrative boundaries. The privacy threat model in multi-domain slicing recognizes that domain operators are not fully trusted by one another, and a malicious aggregator or a compromised domain agent could attempt to infer sensitive operational

parameters from shared model updates. Defenses rooted in differential privacy add calibrated noise to locally computed gradients or to the aggregated global model, providing provable privacy guarantees at the cost of a controlled degradation in policy performance. The choice of privacy budget must balance the sensitivity of slice telemetry — where extremely precise per-user latency figures could reveal individual user behaviors — against the acceptable performance margin for slice assurance. Secure aggregation protocols, which allow the parameter server to compute the sum of encrypted updates without decrypting individual contributions, further strengthen privacy by ensuring that even the aggregating entity cannot inspect single-domain updates. When combined with differential privacy, such schemes yield a robust privacy-preserving substrate that withstands both external eavesdropping and honest-but-curious internal adversaries.

Trusted execution environments offer an alternative hardware-assisted approach, where model aggregation occurs within enclaves that attest to their integrity and confidentiality, preventing the host orchestrator from accessing plaintext parameters. While promising, enclave-based solutions introduce additional dependencies on hardware vendors and complicate deployment across the heterogeneous computing platforms typical of telecom infrastructures, which range from cloud data centers to edge servers and programmable switches. The selection of a privacy mechanism thus involves a multi-faceted evaluation of computational overhead, latency, trust assumptions, and compatibility with the existing security infrastructure of each domain.

Beyond technical protections, federated reinforcement learning architectures must incorporate governance protocols that define data minimization, purpose limitation, and auditability. Data processing agreements among federation participants should clearly specify the nature of information exchanged, the aggregation frequency, the retention period for model updates, and the circumstances under which model inspection is permissible. These agreements operate alongside technical mechanisms to create a layered defense, where even if privacy-preserving computations are imperfectly implemented, contractual and regulatory safeguards deter misuse and provide channels for dispute resolution.

5. Resource Management and Orchestration Strategies

Federated reinforcement learning can be applied to a broad spectrum of resource management tasks in multi-domain slicing, including inter-domain bandwidth reservation, virtual network function placement, edge computing resource allocation, and dynamic spectrum sharing. The local reward functions that guide agent behavior must encode domain-specific performance metrics, such as infrastructure utilization or energy consumption, while also reflecting global service level objectives agreed upon during slice instantiation. Achieving alignment between local incentives and end-to-end performance requires careful design of the credit assignment signal: domain agents must receive feedback that correlates their resource decisions with the satisfaction of cross-domain quality-of-service targets. Multi-objective reinforcement learning frameworks and reward shaping techniques become indispensable tools to balance competing goals such as throughput maximization, latency minimization, fairness among slices, and operational expenditure reduction.

The choice of reinforcement learning algorithm fundamentally shapes the orchestration system’s sample efficiency and adaptability. On-policy methods such as proximal policy optimization offer stable convergence and have been successfully applied to 5G slice assurance tasks [13, 14]. When deployed in a federated setting, on-policy algorithms require careful management of the experience collection phase to ensure that local trajectories reflect

the current global policy while preventing excessive staleness. Off-policy algorithms may improve sample efficiency but raise additional challenges regarding the compatibility of replay buffers with privacy constraints, because storing and potentially sharing experiences could inadvertently leak sensitive information. Hybrid strategies, where domains maintain private replay buffers and synchronize only policy parameters, can mitigate these concerns at the expense of convergence speed.

The temporal dynamics of network traffic and slice demands further complicate learning. Agents must generalize across diurnal patterns, sudden flash crowds, and long-term infrastructure upgrades. Federated meta-learning and continual learning approaches can be integrated to enable rapid adaptation when domains encounter previously unseen traffic regimes, leveraging the collective memory of the federation to bootstrap local policies while avoiding catastrophic forgetting of previously learned resource allocation strategies. The resulting system must orchestrate not only instantaneous resource assignments but also the learning process itself, scheduling training rounds, aggregating updates, and monitoring drift to ensure that the global policy remains robust in the face of evolving operational environments.

6. Robustness, Fairness, and Governance

Multi-domain federated reinforcement learning systems operate in inherently adversarial environments where non-independent and identically distributed data distributions, heterogeneous infrastructure capabilities, and competing economic interests threaten both performance and trustworthiness. Statistical heterogeneity across domains — manifesting as differing traffic profiles, user densities, and slice compositions — can cause naive federated averaging to produce a global policy that is well-tuned only to the dominant statistical regime, leaving minority or resource-constrained domains with degraded slice performance. Personalization techniques, clustered federated learning, and multi-task learning allow the system to maintain a shared representation while adapting policy heads to local conditions, addressing the tension between global coordination and domain-specific optimization [16]. Robustness must also be engineered against adversarial agents that deliberately submit poisoned updates to degrade global policy performance or to bias resource allocation in favor of their own slices. Robust aggregation rules, outlier detection mechanisms, and verifiable computation schemes that allow domains to audit the contributions of peers are essential to sustain cooperative learning in a competitive multi-stakeholder landscape.

Fairness in resource allocation emerges as a particularly complex requirement in multi-domain slicing. Different slices may be owned by competing mobile virtual network operators or enterprises, and the slicing framework must prevent a domain from systematically prioritizing its affiliated slices over those of competitors, even when such preferences are concealed within differential privacy noise or subtle policy actions. Fairness-aware reinforcement learning methods that explicitly incorporate equity constraints into the reward function or policy search process must be adapted to the federated context, where the global notion of fairness must reconcile potentially divergent local fairness criteria [17]. The governance structure of the federation, including the rules for membership, voting on aggregation algorithms, and auditing of compliance with fairness guarantees, therefore becomes as important as the underlying learning algorithm. Multi-stakeholder governance models inspired by internet governance institutions and industry consortia can provide templates for establishing trusted federations where participants have representation in the decision-making processes that shape algorithmic behavior.

7. Deployment and Sustainability Considerations

Translating federated reinforcement learning from simulation testbeds to production multi-domain 5G networks demands attention to practical deployment constraints, including integration with existing management and orchestration platforms, communication overhead, energy footprint, and lifecycle sustainability. The learning infrastructure must interface with 3GPP-compliant management service frameworks, such as the management data analytics service, and leverage standardized information models for slice and resource representation. Automated pipelines for model validation, staged rollout through canary deployments, and rollback capabilities are required to prevent unstable policies from impacting live services. Since federated reinforcement learning involves periodic communication of model updates over wide-area networks, the bandwidth consumption and latency of aggregation rounds must be constrained without sacrificing convergence speed; techniques such as gradient compression, sparse updates, and over-the-air federated learning have been explored in wireless settings [8, 11] and require adaptation to the multi-domain backhaul context.

Sustainability considerations extend beyond energy-efficient radio and compute resource allocation to the broader environmental impact of the learning infrastructure itself. Training reinforcement learning agents, even when federated, consumes significant computational energy across cloud and edge data centers, raising questions about the net carbon benefit of intelligent resource management. Lifecycle assessments that account for model training, inference, and periodic retraining must inform the design of lightweight policy networks and efficient aggregation protocols. The shift toward edge-cloud continuum architectures provides opportunities to co-locate learning workloads with renewable energy sources and to schedule training rounds during periods of low carbon intensity. Moreover, the long-term sustainability of the federated ecosystem depends on the economic viability of participation for smaller infrastructure providers, who may lack the computational resources to train competitive agents. Lightweight federated learning protocols that reduce the client-side computation burden and tiered membership models that allow resource-limited domains to benefit from global knowledge without contributing full-capacity training are necessary to prevent the federation from becoming an exclusive club dominated by large-scale operators.

8. Policy Implications and Future Directions

The deployment of federated reinforcement learning in multi-domain 5G slicing sits at the intersection of telecommunications regulation, data protection law, competition policy, and algorithmic governance. Cross-domain data flows associated with federated model updates may still be subject to international data transfer restrictions, even when raw personal data are not exchanged. Regulators are increasingly scrutinizing whether aggregated model parameters can be considered personal data, particularly when advanced inference attacks can reconstruct sensitive information about individuals or organizational processes [9]. The evolution of regulatory interpretation will directly influence the permissible privacy parameters, aggregation frequencies, and cryptographic safeguards that federation implementers must adopt. Proactive engagement between the research community and regulatory agencies is needed to develop shared technical-legal vocabularies and to avoid regulatory fragmentation that could stall international federations.

Competition policy raises questions about whether federated reinforcement learning consortia could facilitate tacit collusion among participating operators, for instance by implicitly coordinating pricing of slices or jointly limiting capacity investments under the guise of efficient resource allocation. Algorithmic fairness and transparency requirements may compel

federation operators to disclose the decision-making logic of learned policies, a challenging proposition given the inherent opacity of deep neural networks. Explainable reinforcement learning and post-hoc interpretability methods must therefore advance alongside federated architectures to provide stakeholders with the visibility necessary for regulatory compliance and public trust.

Service level agreements in a federated multi-domain environment must evolve beyond static throughput and latency guarantees to encompass algorithmic accountability provisions. Contracts may need to specify acceptable drift in policy performance, maximum tolerable differential privacy noise, and joint liability arrangements for failures caused by cascading policy miscoordination. Dispute resolution mechanisms, possibly leveraging blockchain-based audit trails of model updates and orchestration actions, can provide a verifiable record that aids in post-mortem analysis and liability attribution.

Looking ahead, several research directions merit intensive investigation. The integration of zero-trust security principles with federated reinforcement learning can enable more granular policy enforcement based on continuous verification of participant integrity. Cross-domain digital twins, constructed from synthetic yet statistically faithful network data, may offer safe environments for pre-validating global policies before deployment, reducing the risk of harmful interactions in production. Semantic communication and foundation models for network operations could reduce the dimensionality and sensitivity of shared updates by operating in compact, high-level feature spaces. Ultimately, the maturation of federated reinforcement learning for multi-domain slicing will depend on sustained interdisciplinary collaboration spanning systems engineering, machine learning, law, and economics, forging solutions that are not only technically elegant but also institutionally resilient.

9. Conclusion

The transition to multi-domain 5G network slicing promises unprecedented service customization and infrastructure sharing but simultaneously exposes the limitations of centralized resource management paradigms in privacy-sensitive, competitive ecosystems. Federated reinforcement learning emerges as a transformative architectural framework that enables collaborative policy learning across administrative boundaries without compromising the confidentiality of operational data. This paper has provided a system-level analysis of the key structural, privacy, governance, and deployment dimensions that shape the realization of such federated learning systems. Through detailed examination of aggregation topologies, privacy mechanisms, resource orchestration strategies, fairness and robustness guarantees, and sustainability concerns, a holistic picture of the trade-offs inherent in federated learning for slicing has been constructed. The analysis underscores that technical innovations in differential privacy, secure aggregation, robust personalized learning, and efficient communication must be coupled with robust governance models, transparent accountability mechanisms, and forward-looking policy frameworks to earn the trust of all stakeholders. By synthesizing insights from networking, machine learning, and regulatory scholarship, the paper outlines a research agenda that positions federated reinforcement learning as a cornerstone technology for intelligent, privacy-preserving, and fair multi-domain network orchestration in 5G and beyond.

References

1. Konečný, J., McMahan, H. B., Yu, F. X., Richtárik, P., Suresh, A. T., & Bacon, D. (2016). Federated learning: Strategies for improving communication efficiency. arXiv preprint arXiv:1610.05492.

2. McMahan, B., Moore, E., Ramage, D., Hampson, S., & Agüera y Arcas, B. (2017). Communication-efficient learning of deep networks from decentralized data. In Artificial Intelligence and Statistics (pp. 1273–1282). PMLR.

3. Foukas, X., Patounas, G., Elmokashfi, A., & Marina, M. K. (2017). Network slicing in 5G: Challenges and opportunities. IEEE Internet Computing, 21(1), 20–28.

4. Ordonez-Lucena, J., Ameigeiras, P., Lopez, D., Ramos-Munoz, J. J., Lorca, J., & Folgueira, J. (2017). Network slicing for 5G with SDN/NFV: Concepts, architectures, and challenges. IEEE Communications Magazine, 55(5), 80–87.

5. Zhang, H., Liu, N., Chu, X., Long, K., Aghvami, A.-H., & Leung, V. C. M. (2017). Network slicing based 5G and future mobile networks: A survey. IEEE Communications Surveys & Tutorials, 19(3), 1655–1687.

6. Baranda, J., Mangues-Bafalluy, J., Requena, M., Núñez-Martínez, J., & De La Oliva, A. (2020). Multi-domain orchestration of end-to-end network slices in 5G networks: A reference architecture and prototype. IEEE Transactions on Network and Service Management, 17(4), 2267–2282.

7. Wang, X., Chen, M., Taleb, T., Ksentini, A., & Leung, V. C. M. (2019). Deep reinforcement learning for resource management in network slicing. IEEE Access, 7, 66120–66136.

8. Chen, M., Gündüz, D., Huang, K., Saad, W., Yin, C., & Poor, H. V. (2020). Federated reinforcement learning: Distributed machine learning for next-generation wireless communications. IEEE Vehicular Technology Magazine, 15(3), 61–72.

9. Yin, X., Zhu, Y., & Hu, J. (2021). A comprehensive survey of privacy-preserving federated learning. IEEE Communications Surveys & Tutorials, 23(4), 2400–2442.

10. Lu, Y., Huang, X., Dai, Y., Maharjan, S., & Zhang, Y. (2020). Privacy-preserving federated learning for Internet of Vehicles. IEEE Transactions on Vehicular Technology, 69(12), 15316–15327.

11. Wahab, O. A., Mourad, A., Otrok, H., & Taleb, T. (2021). Federated learning meets 5G: Architectures, enablers, and applications. IEEE Network, 35(3), 46–53.

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

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

14. Schulman, J., Wolski, F., Dhariwal, P., Radford, A., & Klimov, O. (2017). Proximal policy optimization algorithms. arXiv preprint arXiv:1707.06347.

15. Khan, R., Kumar, P., Jayakody, D. N. K., & Liyanage, M. (2020). A survey on security and privacy of 5G technologies: Potential solutions, recent advances, and future directions. IEEE Communications Standards Magazine, 4(4), 58–63.

16. 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.

17. Luo, J., Wu, C., Pan, J., & Liu, Y. (2021). Fairness-aware resource allocation in network slicing: A multi-agent deep reinforcement learning approach. IEEE Transactions on Cognitive Communications and Networking, 7(2), 556–569.

18. Sciancalepore, V., Samdanis, K., Costa-Pérez, X., Banchs, A., Capone, A., & Giannoulakis, I. (2018). Mobile traffic forecasting for 5G network slicing: A deep learning approach based on multi-tenant resource scheduling. IEEE Transactions on Mobile Computing, 17(12), 2903–2916.

19. Buzzi, S., I, C.-L., Klein, T. E., Poor, H. V., Yang, C., & Zappone, A. (2016). A survey of energy-efficient techniques for 5G networks and challenges ahead. IEEE Journal on Selected Areas in Communications, 34(4), 697–709.

20. Marsden, C. T. (2020). Beyond Europe: The internet, regulation, and multistakeholder governance. Telecommunications Policy, 44(4), 101934.