

Knowledge-Augmented Spatial–Temporal Traffic Forecasting for Autonomous Network Management in 6G Systems

Reid Lindgren

Department of Electrical Engineering and Computer Science, University of Missouri,
Columbia, MO, USA.

reidlindgren@missouri.edu

Wayne Ferguson

Department of Computer Science, Colorado State University, Fort Collins, CO, USA.

wayne.ferguson@colostate.edu

Blake Murray

School of Computing, Clemson University, Clemson, SC, USA.

murray1999@clemson.edu

Abstract

The vision of sixth-generation (6G) networks encompasses extreme heterogeneity, near-instantaneous reconfigurability, and autonomous closed-loop management capable of orchestrating resources without human intervention. Central to this ambition is the ability to forecast network traffic with high fidelity across both spatial and temporal dimensions, enabling proactive and predictive resource allocation. While pure data-driven deep learning models have demonstrated remarkable accuracy in spatial-temporal traffic prediction, they often operate as black boxes, neglecting rich external knowledge such as infrastructure semantics, event calendars, social dynamics, and physical environment constraints. Knowledge-augmented forecasting, which fuses structured and unstructured external knowledge with neural spatial-temporal models, has emerged as a promising alternative to overcome these limitations. This paper presents a system-level examination of knowledge-augmented spatial-temporal traffic forecasting for autonomous network management in 6G systems. We articulate an end-to-end architectural vision that integrates heterogeneous knowledge sources, including large language models and knowledge graphs, within frameworks such as the O-RAN RAN Intelligent Controller and ETSI Zero-touch Network and Service Management. The discussion systematically addresses structural trade-offs between centralized and edge-distributed knowledge processing, the computability-latency-accuracy envelope under 6G latency budgets, and the infrastructural requirements for sustainable deployment. An in-depth governance analysis treats robustness against knowledge poisoning and adversarial perturbations, and fairness in resource allocation when predictions incorporate unevenly distributed external knowledge. We further reflect on the sustainability footprint of augmented prediction pipelines and propose strategies for life-cycle-aware design. The paper concludes by outlining future research directions that merge system-level trustworthiness with high-performance predictive autonomy.

Keywords

autonomous network management; 6G systems; spatial-temporal traffic forecasting; knowledge augmentation; zero-touch networks; robustness; fairness; sustainability.

1. Introduction

The transition toward sixth-generation wireless systems heralds a radical re-architecting of communication infrastructures, driven by use cases that span immersive extended reality, holographic telepresence, massive digital twinning, and autonomous vehicular coordination. These scenarios impose stringent demands on user-plane latency, connection density, and spectrum efficiency, while simultaneously expecting networks to self-configure and self-heal at millisecond timescales. Autonomous network management, often conceptualized under the umbrella of zero-touch network and service management, has thus become a cornerstone of 6G design philosophy. Predictive intelligence—particularly the ability to forecast spatio-temporally varying traffic loads—is a foundational enabler of such autonomy, allowing network orchestrators to pre-emptively allocate radio resources, migrate edge workloads, steer beams, and implement energy-saving sleep cycles without waiting for reactive threshold breaches.

Contemporary spatial-temporal traffic forecasting relies heavily on deep learning architectures that jointly model spatial correlations through graph convolutions and temporal dependencies through recurrent or attention-driven mechanisms. Foundational models such as spatio-temporal graph convolutional networks, diffusion convolutional recurrent neural networks, and graph WaveNet have demonstrated impressive predictive accuracy on city-scale and campus-scale datasets [1][2][3]. These models, however, are fundamentally data-driven: they learn statistical patterns from observed traffic time series and static adjacency representations of sensor topology but remain agnostic to the semantic and causal context that generate traffic patterns. Real-world traffic dynamics are shaped by factors such as road closures, weather conditions, scheduled public events, regional power outages, or sudden shifts in mobility caused by social media trends. Capturing such factors demands that prediction systems transcend raw spatio-temporal sequences and incorporate auxiliary knowledge.

Knowledge-augmented spatial-temporal forecasting aims to fill this gap by integrating external structured knowledge graphs, unstructured text streams processed by large language models, and domain-specific ontologies into the predictive workflow. The integration allows models to reason over semantic relationships—for example, inferring that a concert at a stadium increases load on adjacent small cells and backhaul links—even when no direct historical precedent exists in the local traffic logs. A recent example of such augmentation is the TIDES framework, which demonstrates how large language model embeddings can enhance spatial-temporal prediction by encoding world knowledge that complements numerical time series [6]. These techniques open a pathway toward more generalizable and interpretable forecasts, but they also introduce profound system-level challenges that extend well beyond the mere design of neural network layers.

This paper provides a comprehensive system-oriented analysis of knowledge-augmented traffic forecasting as it relates to autonomous network management in 6G environments. Rather than focusing on algorithmic minutiae or incremental accuracy gains, we examine the architectural embedding of augmentation into end-to-end management planes, the structural trade-offs between centralized cloud processing and decentralized edge-native knowledge fusion, the governance requirements for robustness, explainability, and fairness in automated decision loops, and the sustainability implications of computationally heavy knowledge

pipelines. Through a holistic lens, we engage with the policy, deployment, and operational dimensions that will determine whether such augmentation can be safely and equitably instantiated in future public network infrastructures.

2. Background and Motivation

The transition from 5G to 6G represents more than a quantitative increase in peak data rates; it embodies a qualitative shift toward fully autonomous, self-sustaining network fabrics. Architectural visions for 6G, such as those articulated in the broader literature, emphasize the convergence of communication, sensing, and computing, enabled by terahertz bands, reconfigurable intelligent surfaces, and pervasive artificial intelligence embedded throughout the radio access network and core [4]. Within this setting, autonomous network management must handle a complexity explosion arising from ultra-dense small cell deployments, dynamic spectrum sharing, network slicing at fine granularities, and the need to maintain strict service level agreements across diverse verticals.

Zero-touch network and service management provides a structured framework for realizing autonomicity through closed-loop automation, in which monitoring data continuously feeds analytic engines whose insights drive policy-based actuation [5]. The reference architecture specified by the European Telecommunications Standards Institute (ETSI) defines domains where intent-based policies, artificial intelligence analytics, and orchestration functions collaborate without requiring operator intervention for routine decisions. In such a paradigm, the quality of traffic forecasting directly determines the quality and safety of automated actions. A forecast that accurately predicts a surge in demand at a particular edge site enables proactive instantiation of compute instances and caching of popular content before congestion occurs; an inaccurate forecast can trigger unnecessary reconfigurations that waste energy or, worse, fail to avert a degradation in user experience.

Traffic forecasting within these high-stakes loops must capture spatio-temporal variations at multiple scales: millisecond-level fluctuations on a single radio slot, diurnal patterns on a per-cell basis, and regional mobility trends evolving over hours or days. While graph neural network-based models such as spatio-temporal graph convolutional networks [1], diffusion convolutional recurrent neural networks [2], and graph WaveNet [3] have achieved benchmarks on public datasets, they primarily rely on a static or adaptive graph constructed from physical proximity or statistical correlation. Such graphs do not encode why two locations might co-vary due to shared functional roles (e.g., both being train stations served by the same transport line) or due to the causal influence of an external event. Early attempts to incorporate side information, including weather conditions in citywide crowd flow prediction using deep residual networks [3], demonstrated that even simple external features can significantly improve forecast robustness. More recent works have explored the integration of auxiliary spatio-temporal data such as point-of-interest distributions, road network topologies, and event streams learned through sequence models [8].

However, the massive scale and heterogeneity of 6G environments render manual feature engineering of external knowledge untenable. Automated knowledge acquisition and fusion techniques are needed, which has driven interest in leveraging large pre-trained language models and knowledge graphs. These models can ingest unstructured text—ranging from transportation authority announcements to social media posts—and produce dense vector representations that encapsulate semantic context relevant to mobility patterns and communication demand. This marks a shift from simply feeding additional numerical features

toward a genuine knowledge-aware predictive architecture, which then forms the backbone for autonomous decision pipelines.

3. Foundations of Spatial-Temporal Forecasting with Knowledge Augmentation

Spatial-temporal forecasting models designed for traffic networks typically embed sensor locations into a graph where nodes correspond to measurement points—cells, routers, or road segments—and edges encode spatial dependencies. Temporal dynamics are captured by recurrent neural networks, one-dimensional convolutions, or self-attention operators that propagate information across time steps. The combination of spatial graph convolutions and temporal processors has been instantiated in numerous forms, but the core inductive bias remains fixed on the aspect that spatial adjacency is static and known a priori. In real networks, effective spatial relationships are fluid and modulated by the environment, the time of day, and the coherence of user behavioral patterns.

Knowledge augmentation seeks to break the static graph limitation by injecting external semantic information directly into the spatial-temporal learning pipeline. Structured knowledge can be encoded as a knowledge graph where nodes represent geographical zones, network functions, or semantic entities (e.g., “sports venue,” “business district”), and edges indicate relationships such as “is adjacent to,” “served by,” or “hosts event of type.” Knowledge graph embedding techniques transpose these discrete relationships into continuous vector spaces that can be combined with spatial node embeddings [16]. Unstructured knowledge, extracted from natural language using large language models, provides a complementary source of world context. For instance, the TIDES framework demonstrates how DeepSeek, a large language model, can generate contextual embeddings that capture latent factors such as social sentiment, weather forecasts, and institutional calendars, and integrate them with numerical traffic sequences to improve multi-step predictions [6]. This fusion materially enriches the signal available for autonomous decision-making in 6G, where rapid contextual shifts are common.

The fusion process introduces several design questions that transcend pure modelling. Temporal alignment between rapidly sampled network metrics and slowly updated knowledge sources (e.g., hourly weather bulletins, daily event schedules) requires careful synchronization and interpolation strategies. The semantic gap between knowledge graph triples and numerical traffic values must be bridged without forcing artificial discretization of continuous dynamics. Furthermore, the integration must remain computationally viable on the network edge, where inference latency budgets shrink below one millisecond for some 6G control loops. These concerns compel us to consider knowledge augmentation not as a plug-in module but as a first-class architectural component whose design profoundly influences the entire management system.

4. System Architecture for Knowledge-Augmented Traffic Prediction in 6G

We envision a layered architecture that situates knowledge-augmented traffic prediction within the broader autonomous network management fabric. The lowest layer comprises data ingestion and fusion channels that collect real-time telemetry from the radio access network, including time-series metrics of physical resource block utilization, connection counts, queue lengths, and user-plane latency. This layer also ingests external knowledge streams: structured feeds from city information systems, national weather agencies, crowd-sourced event platforms, and unstructured text from news outlets and social media services. A dedicated knowledge curation pipeline transforms these raw streams into a dynamic knowledge graph

and a set of pre-computed language embeddings updated at frequencies appropriate to each source.

Above the ingestion layer, the spatial-temporal feature extraction plane applies graph-based representations of the network topology together with temporal encoders. Knowledge-enhanced node embeddings, derived from the knowledge graph and language model outputs, are attached to sensor nodes as auxiliary attributes, effectively expanding the feature space with semantically meaningful dimensions. This hybrid representation is fed into a prediction engine that can employ attention-based temporal models to learn which aspects of external knowledge are relevant at different future horizons. The output of this engine is a spatio-temporal load forecast that is consumed by a policy reasoning layer.

The policy layer integrates the forecast with intent-driven goals, service level agreements, and resource availability constraints to generate automated actions such as adjusting transmission power profiles, migrating virtual network functions, or initiating content pre-fetching. This closed loop aligns with the O-RAN Alliance architecture for the RAN Intelligent Controller, which already exposes interfaces for hosting machine learning models and enforcement policies within near-real-time and non-real-time timescales [7]. A knowledge-augmented predictor can be deployed as an xApp or rApp in the O-RAN context, receiving enrichment information via a dedicated knowledge service exposed through the Service Management and Orchestration framework. Similarly, the ETSI Zero-touch Network and Service Management reference architecture can accommodate knowledge augmentation within its analytics domain, positioning the knowledge graph as a shared data service accessible to multiple closed loops [5].

An important architectural decision concerns the placement of knowledge processing elements. Centralized data centers can manage the compute-intensive tasks of updating large knowledge graphs and generating language model embeddings; they then distribute the processed embeddings to distributed edge inference nodes that fuse them with local telemetry. This split preserves low-latency inference at the edge while retaining the ability to harness rich external semantics. The trade-off is that the edge nodes rely on the freshness of the centrally computed embeddings, raising questions about synchronization frequency and the impact of stale knowledge on forecast-guided actuations. Alternative architectures fully distribute lightweight knowledge sources, such as localized event calendars managed by regional orchestrators, which avoids cross-region data movement but risks fragmentation of the global knowledge picture.

5. Structural Trade-offs and Infrastructure Considerations

The introduction of knowledge augmentation into the 6G management plane exposes a series of structural trade-offs between predictive accuracy, computational overhead, communication cost, and operational robustness. In a cloud-centralized design, powerful GPU clusters can process global knowledge streams and deliver high-quality embeddings with minimal quantization loss. The downside is the added latency for embedding transport and the vulnerability to backhaul congestion, which can degrade the timeliness of knowledge delivery during the very traffic peaks that most demand accurate predictions. An edge-first design that performs lightweight language model inference on distributed server blades within the access network shortens the knowledge-to-prediction path but must contend with limited computational and memory resources; large language models may need to be heavily distilled or quantised, potentially sacrificing the depth of semantic understanding that makes augmentation attractive.

These competing pressures create a computability-latency-accuracy envelope that must be navigated based on the specific 6G use case. For ultra-reliable low-latency services such as remote surgery or cooperative automated driving, the prediction loop may require worst-case sub-millisecond response, making even compressed large models problematic. In such regimes, knowledge augmentation can take a coarse-grain advisory role, where longer-horizon predictions precompute policy seeds that fast, purely data-driven micro-models can adapt in real time. Conversely, for non-real-time network optimization—such as overnight energy management or capacity planning—the full richness of knowledge-enhanced forecasting can be exploited without tight latency constraints.

Infrastructure heterogeneity further complicates deployment. 6G networks will integrate terrestrial macrocells, street-level small cells, airborne platforms, and satellite backhaul links, each with vastly different compute and storage capabilities. A uniform knowledge augmentation pipeline cannot be imposed across all topologies; the architecture must support graceful degradation where edge nodes automatically scale the complexity of knowledge fusion according to available resources and current network load. This demands standardized interfaces for knowledge services, dynamic model versioning, and life-cycle management of knowledge artifacts, akin to the operation of containerized network functions in cloud-native environments.

Storage and updating of knowledge graphs also impose a non-negligible operational footprint. A knowledge graph covering an entire metropolitan region may contain millions of entities and relationships, requiring continuous updates as events occur. The write throughput and concurrency control needed to keep the graph consistent while serving predictions to hundreds of distributed agents represent a significant system engineering challenge. Implementing incremental graph snapshot distribution, versioned embedding caches, and eventual consistency models can mitigate these issues, but operators must carefully weigh the benefits of richer knowledge against the increased complexity and failure modes of the knowledge infrastructure.

6. Governance, Robustness and Fairness in Autonomous Network Decisions

Incorporating external knowledge into the decision loop of autonomous network management raises critical governance questions that extend beyond traditional performance metrics. The closed-loop autonomy envisioned in zero-touch architectures transfers a degree of control from human operators to artificially intelligent agents. When these agents base actuations on knowledge-augmented forecasts, the provenance, correctness, and auditability of the external knowledge become factors in network accountability. If a flawed event announcement scraped from an unreliable social media source causes a traffic prediction to forecast a fictitious demand surge, the resulting resource misallocation could degrade service for thousands of users. Consequently, organizations must implement robust knowledge vetting pipelines, including source trust scoring, cross-referencing with multiple independent data streams, and human-on-the-loop oversight for high-impact reconfigurations.

Robustness against adversarial manipulation is of paramount concern. Adversarial perturbations added to input traffic observations can deceive spatial-temporal models, but knowledge augmentation introduces additional attack surfaces. Malicious actors might inject poisoned knowledge triples into public knowledge bases or generate misleading textual content designed to bias language model embeddings. The resulting knowledge-embedded features could propagate through the predictor and cause systematic underestimation or overestimation of traffic in targeted geographic areas, enabling denial-of-service or resource

squatting attacks. Defending against such threats requires a combination of techniques, including adversarial training that encompasses knowledge input modalities, certified robustness methods adapted to graph neural network architectures [9], and continuous monitoring of prediction residuals to detect anomalous deviations that may indicate knowledge poisoning.

Fairness emerges as a cross-cutting governance pillar. Knowledge augmentation can inadvertently encode societal biases that are reflected in external data sources. For example, event data and social media activity may be richer and more timely for affluent urban districts with high smartphone penetration and active digital presence, while rural or economically disadvantaged areas may generate sparser external signals. A knowledge-augmented predictor that relies disproportionately on such signals could systematically under-predict traffic in underserved regions, leading a zero-touch orchestrator to allocate fewer resources to those areas and thereby reinforcing a cycle of inadequate service quality. Mitigating this requires fairness-aware optimization objectives during both the knowledge fusion and the resource allocation stages, as well as explicit auditing of prediction disparities across demographic and geographic categories. Drawing on fairness frameworks in machine learning, network operators can impose constraints that equalize the allocation of excess capacity across slices in a manner sensitive to the distribution of epistemic uncertainty stemming from knowledge sparsity [10][12].

Explainability and transparency are essential to both trust and compliance with emerging regulatory frameworks. When an automated action—such as dynamically re-prioritizing a slice’s resource warrants—is justified by a knowledge-augmented forecast, the decision rationale must be traceable. Operators and regulators may require human-intelligible explanations that identify which external knowledge elements contributed most to the forecast. This can be achieved via attention visualizations, Shapley value decompositions extended to knowledge tokens, or surrogate modelling approaches that approximate the augmented predictor with an interpretable counterpart. Establishing such explainability mechanisms is a pre-condition for the large-scale adoption of closed-loop autonomy in public communication infrastructures and must be embedded into the architecture by design, not retrofitted.

Policy and regulatory dimensions further frame governance. National telecom regulators may mandate that any external data used for essential network functions comply with data privacy laws and be subject to algorithmic impact assessments. Cross-border data flows involved in global knowledge graph sharing raise jurisdictional tensions that require federated knowledge management solutions, where raw data remains within sovereign boundaries while aggregated embeddings are exchanged [14]. Standardization bodies such as the International Telecommunication Union and the 3rd Generation Partnership Project will need to develop guidelines for knowledge-augmented network intelligence, specifying acceptable data sources, audit trail formats, and conformance testing procedures to ensure interoperability and public interest protection.

7. Deployment Strategies and Sustainability Implications

Incremental deployment of knowledge-augmented forecasting will most likely begin within existing 5G-Advanced and Open RAN environments as an overlay intelligence service before being natively designed into 6G greenfield deployments. Network operators can initially deploy knowledge supplements for non-real-time use cases—such as overnight energy optimization and next-day capacity planning—where the operational risk of inaccurate external knowledge is lower and the cycle time permits human vetting. As confidence in

knowledge quality and model robustness grows, the augmentation can be progressively integrated into near-real-time and real-time loops, guided by empirical guardrails that monitor forecast drift and automatically fall back to purely data-driven models when external knowledge sources exhibit anomalies.

Sustainability implications of knowledge augmentation must be examined through a life-cycle lens. On the positive side, more accurate traffic forecasting can significantly reduce the energy consumption of the radio access network by enabling precise and confident sleep mode scheduling, reducing unnecessary signalling and transmission, and optimizing the placement of computing workloads to minimize data movement. Studies on the energy efficiency improvements achievable through intelligent resource allocation in 6G consistently point to advanced prediction as a critical lever [11]. When knowledge augmentation reduces forecast error margins by integration of causal context, the energy savings from avoided over-provisioning can be substantial.

However, the knowledge-processing pipeline itself carries an environmental cost. Training large language models and maintaining continuously updated knowledge graphs require considerable computational resources and associated carbon emissions. The execution of knowledge-enhanced inference at scale across thousands of edge nodes adds operational power draw that must be netted against the efficiency gains. A sustainable deployment thus demands rigorous carbon-aware orchestration of knowledge workloads, scheduling heavy embedding refreshes during periods of high renewable energy availability and employing knowledge distillation to shrink model footprints. Furthermore, the concept of a federated knowledge ecosystem, in which operators collaboratively contribute anonymized gradients or embedding updates without centralizing raw data, can amortize the sustainability penalty across stakeholders while enhancing knowledge freshness for all participants [14][18].

Digital twin platforms offer a compelling vehicle for safe, sustainable pre-deployment validation. A high-fidelity digital replica of the target 6G network, enriched with simulated knowledge streams, allows operators to stress-test knowledge-augmented forecasting pipelines under adversarial scenarios, extreme event sequences, and rare failure modes that might never appear in historical logs [13]. By incorporating energy models and carbon intensity signals directly into the twin, operators can optimize knowledge update policies and resource configurations for minimal operational emissions before any hardware is deployed. Such practices align with the broader industry trajectory toward network life-cycle assessment and will increasingly be mandated by sustainability-linked financing and government net-zero targets.

Eventually, the governance and sustainability aspects converge on the need for accountability frameworks that track the resource and carbon footprint of autonomous decisions. A management agent that decides to activate an additional knowledge source for a forecast improvement implicitly incurs a computational cost and its associated emissions. Future policy instruments may require operators to justify that the expected service quality uplift justifies the environmental impact, particularly for routine operations. Embedding carbon ledgering into the autonomous management plane and linking it with the knowledge curation subsystem will provide the necessary transparency to navigate this socio-technical trade-off.

8. Conclusion

This paper has presented a system-level perspective on knowledge-augmented spatial-temporal traffic forecasting as an enabling component of autonomous network management in

6G systems. By moving beyond purely data-driven models, knowledge augmentation holds the promise of more robust, generalizable, and context-aware predictions that can underpin true zero-touch orchestration. However, reaping these benefits demands careful attention to architectural integration, structural trade-offs, governance, fairness, and sustainability, all of which interact non-trivially. The embedding of augmented forecasting within standardized frameworks such as the O-RAN RIC and ETSI ZSM provides a pathway for incremental adoption, while the design of centralized-distributed knowledge processing fabrics determines whether the approach can satisfy the extreme latency and reliability requirements of future verticals. Governance instruments must evolve to address knowledge poisoning, democratic access to high-quality external data, and algorithmic accountability, while sustainability mandates will shape the operational envelope of knowledge-intensive prediction pipelines. As research communities and industry consortia further develop the TIDES-like frameworks and knowledge fusion methods, the holistic considerations articulated here must remain at the forefront to ensure that the technology serves society equitably, robustly, and within planetary boundaries.

References

1. Yu, B., Yin, H., & Zhu, Z. (2018). Spatio-temporal graph convolutional networks: A deep learning framework for traffic forecasting. *Proceedings of the 27th International Joint Conference on Artificial Intelligence (IJCAI-18)*, 3634–3640.
2. Li, Y., Yu, R., Shahabi, C., & Liu, Y. (2018). Diffusion convolutional recurrent neural network: Data-driven traffic forecasting. *International Conference on Learning Representations (ICLR)*.
3. Zhang, J., Zheng, Y., & Qi, D. (2017). Deep spatio-temporal residual networks for citywide crowd flows prediction. *Proceedings of the AAAI Conference on Artificial Intelligence*, 31(1), 1655–1661.
4. 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.
5. ETSI. (2019). Zero-touch network and service management (ZSM); Reference architecture (ETSI GS ZSM 002 V1.1.1). European Telecommunications Standards Institute.
6. Zhang, C., Zhang, H., Qiao, J., Li, Z., & Alouini, M. S. (2025). TIDES: Traffic Intelligence with DeepSeek Enhanced Spatial Temporal Prediction. *IEEE Journal on Selected Areas in Communications*.
7. O-RAN Alliance. (2020). O-RAN: Towards an open and smart RAN. O-RAN Alliance White Paper.
8. Liao, B., Zhang, J., Wu, C., McIlwraith, D., Chen, T., Yang, S., ... & Guo, Y. (2018). Deep sequence learning with auxiliary information for traffic prediction. *Proceedings of the 24th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining*, 834–843.
9. Zügner, D., & Günnemann, S. (2019). Adversarial attacks on graph neural networks via meta learning. *International Conference on Learning Representations (ICLR)*.
10. Eisen, M., Rashid, M. M., & Papachristodoulou, A. (2019). Learning fair resource allocation policies in wireless networks. *IEEE Transactions on Cognitive Communications and Networking*, 5(3), 787–798.

11. Alsabah, M., Naser, M. A., Mahmmmod, B. M., Abdulhussain, S. H., Eissa, M. R., Al-Baidhani, A., ... & Ahmed, H. (2021). 6G wireless communications networks: A comprehensive survey. *IEEE Access*, 9, 148191–148218.
12. Mehrabi, N., Morstatter, F., Saxena, N., Lerman, K., & Galstyan, A. (2021). A survey on bias and fairness in machine learning. *ACM Computing Surveys*, 54(6), 1–35.
13. Shen, X., Fantacci, R., & Hanzo, L. (2022). Digital twin for 6G: Architecture, use cases, and open issues. *IEEE Network*, 36(3), 10–17.
14. Yang, Q., Liu, Y., Chen, T., & Tong, Y. (2019). Federated machine learning: Concept and applications. *ACM Transactions on Intelligent Systems and Technology*, 10(2), 1–19.
15. Strubell, E., Ganesh, A., & McCallum, A. (2019). Energy and policy considerations for deep learning in NLP. *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, 3645–3650.
16. Wang, Q., Mao, Z., Wang, B., & Guo, L. (2017). Knowledge graph embedding: A survey of approaches and applications. *IEEE Transactions on Knowledge and Data Engineering*, 29(12), 2724–2743.
17. Zhou, J., Cui, G., Zhang, Z., Yang, C., Liu, Z., Wang, L., ... & Sun, M. (2020). Graph neural networks: A review of methods and applications. *AI Open*, 1, 57–81.
18. Khan, L. U., Saad, W., Han, Z., & Hossain, E. (2022). Energy efficient federated learning over wireless communication networks. *IEEE Transactions on Green Communications and Networking*, 6(3), 1738–1749.
19. Samek, W., Montavon, G., Lapuschkin, S., Anders, C. J., & Müller, K.-R. (2021). Explaining deep neural networks and beyond: A review of methods and applications. *Proceedings of the IEEE*, 109(3), 247–278.
20. Zhao, J., Li, Q., Gong, Y., & Zhang, Y. (2021). AI-driven zero-touch network and service management in 6G: A vision. *IEEE Network*, 35(3), 98–105.