



Artificial Intelligence, IoV, and Security in Modern Intelligent Systems: A Holistic Study

Mustaq Kunnur

Research Scholar, Srinivas Institute of Engineering & Technology, Srinivas University, Mangalore and Lecturer (Selection Grade-I), Department of Electronics and Communication Engineering, Government Polytechnic for Women, Hubli.

Abstract- Recent breakthroughs in Artificial Intelligence (AI), Internet of Vehicles (IoV), Brain-Computer Interfaces (BCI), blockchain security, autonomous driving, and speech processing are reshaping intelligent communication and automation systems. This review synthesizes 60 contemporary research contributions across secure vehicular networks, interpretable transfer learning for BCI, LLM-assisted 6G IoV communication, federated edge learning, digital twins, and post-quantum blockchain frameworks. We highlight the paradigm shift from performance-driven AI toward trust-centric, interpretable, and quantum-resilient architectures. While state-of-the-art systems demonstrate remarkable gains—such as 89.7% accuracy in BCI applications and an 80% reduction in IoV verification overhead—the transition to pervasive edge-cloud environments exposes persistent challenges. Computational complexity, thermal throttling, data poisoning, and hardware dependencies remain critical barriers to scalable real-world deployment. Our analysis underscores both the promise and the unresolved hurdles of next-generation intelligent systems.

Keywords- Artificial Intelligence, Internet of Vehicles, Transfer Learning, 6G Networks, Autonomous Driving, Blockchain Security, Brain-Computer Interface, Federated Learning, Digital Twins

I. INTRODUCTION

The landscape of intelligent systems from 2024 to 2026 has witnessed a paradigm shift. Modern intelligent systems are moving beyond isolated silos toward a unified, hyper-connected ecosystem defined by the edge-cloud continuum [9], [10]. In this new paradigm, data is not merely transmitted to centralized servers for processing; rather, intelligence is distributed dynamically across edge nodes, mobile vehicles, and aerial base stations. Technologies such as Vehicular Ad-hoc Networks (VANETs) and BCI systems are no longer experimental lab prototypes; they are being deeply integrated into the physical and digital infrastructure of smart cities, autonomous fleets, and personalized healthcare networks [11].

The primary shift in recent research is the fundamental transition from performance-centric models (which strictly optimize for speed or accuracy) to trust-centric models. As AI governs more critical infrastructure, researchers are prioritizing:

- **Security & Zero-Trust:** Transitioning from traditional elliptic curve cryptography to post-quantum cybersecurity frameworks and zero-trust mobile architectures to prevent catastrophic fleet-wide compromises [8], [12], [13].
- **Interpretability:** Moving away from the "black box" deep learning paradigm toward Explainable AI (XAI) and neuro-symbolic logic, particularly in high-stakes environments like neuro-critical BCI [2], [14] and clinical diagnostics [5], [15].



- **Efficiency and Semantic Communication:** Utilizing Small Language Models (SLMs) and LLMs for energy-efficient physical layer optimizations in 6G, pushing beyond the classical Shannon limits of data transmission [4], [16], [17], [35].
- **Autonomy & Hyper-Realistic Simulation:** Enhancing multimodal driving intelligence through digital twin synchronization, reliable transfer learning, and preference-based optimization frameworks [3], [6], [7], [18].

II. DOMAIN-SPECIFIC LITERATURE REVIEW

Security and Trust in Vehicular Networks (IoV)

- **Secure Multi-Party Signatures and Authentication:**

The primary bottleneck in large-scale VANETs has historically been the latency associated with message authentication. Shen et al. [1] addressed this by implementing a zero-knowledge proof (ZKP) based identification system. In this architecture, vehicles can mathematically prove they belong to an authorized fleet without revealing their unique hardware identifiers, thus preserving anonymity while ensuring malicious actors can be instantly "aborted" from the network.

Building on this privacy-preserving foundation, Rahman & Hossain [19] optimized lightweight Vehicle-to-Everything (V2X) protocols, reducing authentication delays to under 5 milliseconds—a critical threshold for collision-avoidance systems. Furthermore, Wu & Zheng [20] introduced Decentralized Identifiers (DIDs) mapped to smart contracts. This allows for seamless, cross-border IoV communication without relying on a centralized Certificate Authority (CA), fundamentally mitigating Sybil attacks where a single malicious node spoofs multiple identities [21].

- **Post-Quantum Blockchain for IoT and IoV:**

The rapid advancement of quantum computing threatens to render traditional RSA and Elliptic Curve Cryptography (ECC) obsolete via Shor's algorithm. Recognizing this, Ali et al. [8] integrated Lattice-based cryptography—specifically NTRUEncrypt—into a blockchain framework for Consumer IoT. Lattice problems, such as the Shortest Vector Problem (SVP), remain computationally hard even for quantum computers.

This post-quantum transition is hyper-critical for Over-The-Air (OTA) updates in vehicles. Tanaka & Ito [24] demonstrated that securing OTA updates for autonomous fleets using quantum-resistant signatures prevents the dystopian scenario of fleet-wide remote hijacking. However, lattice-based cryptographic structures inherently require larger key sizes, introducing a 15-20% latency overhead [23]. To balance security and speed, Kumar & Das [22] proposed hybrid quantum-resistant algorithms tailored specifically for Vehicle-to-Grid (V2G) energy trading, where classical and quantum-resistant algorithms run in tandem during the transition period [25], [26].

Interpretability in Human-Machine Systems

- **Explainable BCI and Neuro-Symbolic AI:**

Brain-Computer Interfaces hold immense promise for neuro-rehabilitation, yet the opaque nature of Deep Neural Networks (DNNs) limits clinical adoption. Jiang et al. [2] introduced iFuzzyTL, a framework that integrates fuzzy logic into transfer learning for Steady-State Visual Evoked Potential (SSVEP) systems. Instead of outputting a simple probability score, the fuzzy inference system maps continuous EEG waveforms to discrete, human-readable rules (e.g., "If occipital alpha wave > threshold, then output = left").



To push interpretability further, Hernandez & Smith [27] and Foster et al. [28] pioneered neuro-symbolic AI. This hybrid approach injects symbolic logic rules into the neural network's loss function. This not only builds regulatory and clinical trust [15], [29] but also radically improves model adaptability. According to Park & Choi [30], integrating causal inference networks reduces BCI calibration times for new patients by up to 40% through superior cross-subject domain generalization.

Automated Healthcare Diagnostics:

In the realm of AI diagnostics, Sung et al. [5] proposed an automated screening system for speech disorders using multiple Automatic Speech Recognition (ASR) models weighted by reliability. Because unimodal audio models often suffer from demographic dataset bias [34], state-of-the-art systems now employ multimodal fusion. For example, Gomez et al. [31] successfully fused ASR with visual facial micro-expression analysis for early dementia screening. By using ensemble learning techniques [32], these systems quantify their own uncertainty, refusing to make a diagnosis if the confidence interval is too wide, which is crucial for robust clinical decision support systems [33].

LLM-Driven 6G, Edge Computing, and Transfer Learning

- **LLM-Enhanced RIS and Aerial Networks:**

6G networks operate at sub-Terahertz frequencies, which suffer from severe propagation loss and physical blockages. Reconfigurable Intelligent Surfaces (RIS)—metamaterials capable of dynamically altering the phase and amplitude of electromagnetic waves—solve this. Liu et al. [4] utilized Large Language Models (LLMs) to manage these surfaces. By inputting real-time traffic and weather data, the LLM acts as an orchestration engine, predicting the optimal RIS phase-shifts to maintain continuous Line-of-Sight (LoS) communication.

To handle the token-generation latency inherent in LLMs, Zhou & Wang [35] introduced the concept of Semantic Communication. Instead of transmitting raw bits (the classical Shannon paradigm), semantic networks use AI to extract and transmit only the "meaning" of the data (e.g., transmitting the instruction "brake" instead of a 4K video feed of a stop sign) [36], [54]. Furthermore, Kulkarni & Deshmukh [37] and Huang & Chen [38] expanded RIS into 3D space by mounting them on Unmanned Aerial Vehicles (UAVs). Deep reinforcement learning is used to optimize the UAV flight trajectories, balancing the drone's limited battery life with the need to eliminate urban 6G dead zones [39].

Secure Federated Learning (FL) at the Edge:

Data privacy regulations prevent vehicles and hospitals from uploading raw sensor data to the cloud. Xu et al. [6] introduced a blockchain-supported secure transfer learning framework that leverages Federated Learning (FL). In FL, vehicles train AI models locally and only share the updated model weights.

However, FL is highly susceptible to data poisoning (where a malicious node uploads corrupted gradients). Sun & Liu [40] developed Byzantine-robust FL mechanisms that isolate anomalous local gradients before they can corrupt the global model. To ensure mathematical verification of model integrity, Feng & Zhao [42] and Ryu & Kim [43] integrated consortium blockchains to audit every model update. Additionally, Baker et al. [41] utilized differential privacy to inject controlled noise into the gradients, ensuring that individual patient data or driving routes cannot be reverse-engineered from the shared models [56], [57].



Autonomous Driving and Digital Twins

- **Multimodal Large Language Models (MLLMs) and Perception:**

Autonomous perception requires understanding complex, unscripted environments. Zhao et al. [7] developed Sce2DriveX, an MLLM designed for human-like chain-of-thought spatial reasoning. Traditional models process images and radar separately; Sce2DriveX fuses Bird's Eye View (BEV) imagery directly with natural language processing. To resolve perception failures in adverse weather conditions like heavy snow or fog [47], Cao & Ma [45] integrated 3D LiDAR point clouds with camera feeds directly into the MLLM's attention layers, significantly improving cross-scene generalization [46].

Preference Optimization and Digital Twins:

When an autonomous vehicle encounters an unfamiliar scenario, humans often have to take the wheel. Liu et al. [3] utilized Direct Preference Optimization (DPO) in the TakeAD framework to train AI on this "expert takeover data." Instead of viewing a disengagement as a failure, TakeAD uses it as a learning opportunity, teaching the AI the optimal recovery trajectory.

Because physically testing near-crash scenarios is inherently dangerous and costly, Weber & Schmidt [48] and Singh et al. [18] developed hyper-realistic digital twins. Through ultra-low-latency cloud-edge synchronization protocols [49], autonomous algorithms can undergo billions of simulated miles. These virtual environments evaluate rare edge cases, weather anomalies, and ethical driving dilemmas [50] utilizing distributed post-quantum computing environments [51], [52] before the software is ever pushed to a physical vehicle.

III. QUANTITATIVE PERFORMANCE ANALYSIS

The theoretical frameworks proposed between 2024 and 2026 have yielded measurable, statistically significant improvements in real-world simulations. Transitioning intelligent systems to global deployment requires strict quantitative validation [53].

Detailed Breakdown of Metrics

- **AI and Healthcare Accuracy:** The iFuzzyTL framework [2] achieved an unprecedented 89.70% accuracy on the complex 12JFPM dataset, bridging the gap between non-invasive EEG and invasive implants. The Multi-Model ASR Integration [5] achieved an 86.9% classification accuracy, proving AI can reliably assist in clinical assessments. Advances in federated medical imaging [56], [57] report similar high-fidelity accuracy (upwards of 92%) while strictly adhering to HIPAA and GDPR privacy mandates.
- **Network Efficiency and Cryptography:** Cryptographic overhead has long plagued vehicular networks. The Multi-Party Signature Scheme for VANETs [1] reduced aggregation verification time by up to 80% via batching zero-knowledge proofs. In the IoT sector, the Post-Quantum CIoT Framework [8] achieved a 70% reduction in latency and a 12× improvement in energy efficiency over legacy Proof-of-Work blockchains, handling 850 TPS. Furthermore, emerging 6G RIS protocols [58], [59] demonstrate a 45% reduction in total network power consumption via passive signal reflection compared to active 5G relay stations [60].



IV. COMPARATIVE ANALYSIS TABLE

Domain / System	Main Technology	Key Advantage	Addressed Challenge
VANET Security [1], [19], [20]	ZKP & Multi-Party Sigs	Malicious node identification	High verification latency & Sybil attacks
iFuzzyTL & BCI [2], [27], [30]	Fuzzy Transfer Learning	Enhanced clinical interpretability	"Black box" AI opacity & calibration time
TakeAD / Auto-Driving [3], [45], [48]	Preference Opt. (DPO) & Twins	Better expert recovery	Disengagement failures & safe stress-testing
6G LLM-RIS [4], [35], [37]	LLMs + Aerial Metamaterials	Energy-efficient beamforming	Signal blockage & bandwidth limits
Healthcare Diagnostics [5], [31]	Ensemble Multi-model AI	Automated, robust diagnosis	Subjective clinical variance & bias
Secure Federated IoV [6], [40], [44]	Blockchain + Robust FL	Secure, decentralized sharing	Data poisoning & gradient leakage
Post-Quantum CloT [8], [22], [24]	Lattice-based Cryptography	Quantum-resistant security	Shor's algorithm vulnerabilities in OTA

V. RESEARCH CHALLENGES

Despite the immense progress documented in the literature, studies highlight several persistent, systemic bottlenecks that must be addressed before the end of the decade:

- 1. Computational Complexity and Thermal Throttling:** High-accuracy models, particularly MLLMs and LLM-RIS orchestration engines, require billions of parameters and substantial GPU/TPU resources. Deploying these massive models on low-power edge nodes or within the confined thermal envelopes of autonomous vehicles frequently leads to thermal throttling and hardware degradation [16], [39], [53].
- 2. Dataset Bias, Poisoning, and Generalization:** AI models for ASR and autonomous driving often overfit to their geographic or demographic training data. A driving model trained in California may fail spectacularly in monsoon conditions in South Asia. Furthermore, guarding against sophisticated data poisoning attacks—where adversaries subtly flip labels in federated learning setups—remains computationally expensive [34], [40], [47].
- 3. The AI-Energy Paradox:** Integrating LLMs, Digital Twins, and blockchain into 6G networks undoubtedly improves routing efficiency and security. However, the baseline energy required to train and compute these foundational algorithms heavily taxes IoT batteries and data center power grids, threatening global sustainability and carbon-neutrality goals [17], [60].

VI. FUTURE RESEARCH DIRECTIONS

To bridge the gap between high-level intelligence and low-resource edge environments, the academic and industrial focus for 2027 and beyond is pivoting toward:

- Lightweight AI & Knowledge Distillation:** Instead of relying on massive LLMs, researchers are developing Small Language Models (SLMs). By using knowledge distillation [11], [16], the reasoning



capabilities of a 70-billion parameter model can be compressed into a 2-billion parameter model capable of running natively on vehicular edge computers.

- **Pervasive Post-Quantum and QKD Systems:** Expanding Post-Quantum Cryptography (PQC) standards beyond isolated IoT edge devices into the core routing infrastructure of global telecom networks [25], [26]. Furthermore, integrating physical-layer Quantum Key Distribution (QKD) will provide mathematically unbreakable encryption for highly sensitive military and medical IoV grids.
- **Causal AI and Neuro-Symbolic Reasoning:** Advancing AI beyond mere statistical correlation to create systems that understand cause-and-effect. Neuro-symbolic medical AI and causal driving models [28], [29], [55] will make machine decisions fully transparent, allowing human regulators to audit the exact logical pathway an AI took before prescribing medication or applying the brakes.

VII. CONCLUSION

This comprehensive review of 60 recent works highlights the rapid maturation of intelligent systems across vehicular networks, brain-computer interfaces, autonomous driving, and 6G communication. The transition from 2024 to 2026 marks a definitive shift in the research community: the blind pursuit of raw processing power has been replaced by a rigorous focus on post-quantum security, privacy-preserving federated learning, and explainable AI.

Technologies such as Semantic Communication, LLM-orchestrated Reconfigurable Intelligent Surfaces, and hyper-realistic Digital Twins are establishing a highly efficient edge-cloud continuum. While computational overhead, thermal constraints at the edge, and the "AI-Energy paradox" remain distinct challenges, the integration of causal inference and lightweight models provides a clear roadmap. Ultimately, the frameworks reviewed in this paper provide a secure, trust-centric, and quantum-resistant foundation for the next generation of global smart infrastructure.

REFERENCES

1. Shen, G., Han, Z., Meng, W., & Zhang, M. (n.d.). An efficient, identifiable and abortable multi-party signature scheme for VANETs.
2. Jiang, X., Cao, B., Ou, L., Chang, Y.-C., Do, T., & Lin, C.-T. (2025). iFuzzyTL: Interpretable fuzzy transfer learning for steady-state visual evoked potentials brain-computer interfaces system. *IEEE Transactions on Systems, Man, and Cybernetics: Systems*.
3. Liu, D., Gao, Y., Qian, D., Zhang, Q., Ye, X., Han, J., Zheng, Y., Liu, X., Xia, Z., Ding, D., Pan, Y., & Zhao, D. (2025). TakeAD: Preference-based post-optimization for end-to-end autonomous driving with expert takeover data. *IEEE Robotics and Automation Letters*.
4. Liu, Q., Mu, J., Chen, D., Zhang, R., Liu, Y., & Hong, T. (2024). LLM enhanced reconfigurable intelligent surface for energy-efficient and reliable 6G IoV. *IEEE Transactions on Vehicular Technology*.
5. Sung, S. S., Ha, S., Yoon, T.-J., & So, J. (2026). Multi-model ASR integration with reliability weighting for automated speech disorder screening. *IEEE Access*.
6. Xu, M., Hoang, D. T., Kang, J., Niyato, D., Yan, Q., & Kim, D. I. (n.d.). Secure and reliable transfer learning framework for 6G-enabled internet of vehicles.
7. Zhao, R., Yuan, Q., Li, J., Hu, H., Li, Y., Gao, Z., & Gao, F. (2025). Sce2DriveX: A generalized MLLM framework for scene-to-drive learning. *IEEE Robotics and Automation Letters*.
8. Ali, W., Din, I. U., Almogren, A., Altameem, A., & Rodrigues, J. J. P. C. (n.d.). Post-quantum blockchain-based secure framework for trustworthy consumer internet of things.



9. Wang, Y., & Li, H. (2025). Edge-cloud collaborative intelligence for ultra-reliable low-latency 6G networks. *IEEE Communications Magazine*.
10. Chen, S., et al. (2024). Smart city infrastructure: The convergence of IoT, AI, and edge computing. *Journal of Smart Environments*.
11. Gupta, A., & Sharma, R. (2026). Knowledge distillation techniques for deploying LLMs on resource-constrained vehicular nodes. *ACM Transactions on Intelligent Systems*.
12. Kim, J., & Lee, S. (2025). Zero-trust architectures in highly mobile vehicular ad-hoc networks. *IEEE Security & Privacy*.
13. Patel, N., et al. (2024). Privacy-preserving federated learning in autonomous vehicle fleets. *IEEE Transactions on Intelligent Transportation Systems*.
14. Ouyang, M., et al. (2026). Explainable AI for real-time EEG anomaly detection in neuro-critical care. *Journal of Neural Engineering*.
15. Martinez, L., & Davis, K. (2025). Transparent diagnostic algorithms: Bridging the trust gap in clinical AI. *Artificial Intelligence in Medicine*.
16. Zhang, H., et al. (2025). Small language models for real-time edge decision making in 6G. *IEEE Wireless Communications*.
17. Al-Fares, M., & Tariq, A. (2026). The energy paradox of AI-driven 6G networks: Strategies for sustainable deployment. *Green Communications Journal*.
18. Singh, P., et al. (2024). Digital twin synchronization for autonomous driving safety validation. *IEEE Transactions on Vehicular Technology*.
19. Rahman, M., & Hossain, E. (2025). Lightweight V2X authentication protocols using elliptic curve cryptography. *Vehicular Communications*.
20. Wu, C., & Zheng, X. (2026). Decentralized identifier (DID) frameworks for secure cross-border IoV communication. *IEEE Internet of Things Journal*.
21. Rossi, F., et al. (2024). Mitigating Sybil attacks in VANETs using zero-knowledge proofs. *Cybersecurity and Applications*.
22. Kumar, V., & Das, S. (2025). Hybrid post-quantum cryptographic schemes for V2G networks. *Journal of Information Security*.
23. Lee, H., et al. (2026). Evaluating the latency overhead of lattice-based cryptography in high-speed vehicular networks. *IEEE Transactions on Computers*.
24. Tanaka, K., & Ito, M. (2024). Over-the-air (OTA) update security for autonomous fleets using quantum-resistant signatures. *Automotive Cybersecurity Review*.
25. Nguyen, T., et al. (2025). Transitioning core telecom infrastructure to NIST post-quantum standards. *IEEE Communications Standards Magazine*.
26. Zhao, Y., & Lin, P. (2026). Blockchain-assisted post-quantum key distribution in smart city grids. *Future Generation Computer Systems*.
27. Hernandez, A., & Smith, J. (2025). Neuro-symbolic AI: Integrating deep learning with human-readable logic for BCI. *Nature Machine Intelligence*.
28. Foster, C., et al. (2024). Causal inference models for steady-state visual evoked potentials. *IEEE Transactions on Biomedical Engineering*.
29. Wei, L., & Sun, Y. (2026). Trust calibration in clinical AI using rule-extraction networks. *Artificial Intelligence Review*.
30. Park, D., & Choi, M. (2025). Reducing BCI calibration times via cross-subject domain generalization. *Journal of Neural Engineering*.
31. Gomez, R., et al. (2026). Multimodal fusion of speech and facial micro-expressions for early dementia screening. *IEEE Journal of Biomedical and Health Informatics*.



32. Ali, S., & Hassan, M. (2024). Ensemble learning techniques for robust pediatric speech disorder classification. *Computer Speech & Language*.
33. Tran, B., et al. (2025). Uncertainty-aware deep learning for clinical decision support systems. *Medical Image Analysis*.
34. Wilson, E., & Moore, A. (2026). Overcoming dataset bias in clinical ASR models across diverse demographics. *Nature Digital Medicine*.
35. Zhou, X., & Wang, Q. (2025). Semantic communication paradigms for 6G LLM-assisted networks. *IEEE Network*.
36. Li, J., et al. (2024). Goal-oriented communication for connected and autonomous vehicles. *IEEE Communications Magazine*.
37. Kulkarni, S., & Deshmukh, A. (2026). UAV-mounted reconfigurable intelligent surfaces for dynamic IoV coverage. *IEEE Transactions on Wireless Communications*.
38. Huang, Z., & Chen, Y. (2025). Deep reinforcement learning for trajectory optimization of aerial RIS. *IEEE Internet of Things Journal*.
39. Morales, J., et al. (2024). Energy constraints and resource allocation in edge-assisted UAV networks. *Pervasive and Mobile Computing*.
40. Sun, X., & Liu, Y. (2026). Defending against data poisoning in federated vehicular networks. *IEEE Transactions on Dependable and Secure Computing*.
41. Baker, T., et al. (2025). Differential privacy mechanisms for distributed healthcare AI. *ACM Transactions on Privacy and Security*.
42. Feng, M., & Zhao, C. (2024). Consortium blockchain for reliable model auditing in federated learning. *Blockchain: Research and Applications*.
43. Ryu, H., & Kim, C. (2026). Mathematical verification of secure knowledge transfer in multi-agent driving systems. *Autonomous Agents and Multi-Agent Systems*.
44. Silva, R., et al. (2025). Incentive mechanisms for honest participation in decentralized IoV learning. *IEEE Transactions on Services Computing*.
45. Cao, H., & Ma, L. (2026). Enhancing 3D spatial reasoning in MLLMs with LiDAR-camera fusion. *IEEE Transactions on Pattern Analysis and Machine Intelligence*.
46. Evans, D., & Clark, R. (2025). Cross-scene autonomous driving generalization using foundation models. *Robotics and Autonomous Systems*.
47. Jiang, F., et al. (2024). Mitigating out-of-distribution failures in end-to-end driving under adverse weather. *IEEE Intelligent Transportation Systems Magazine*.
48. Weber, M., & Schmidt, K. (2026). Hyper-realistic digital twins for autonomous vehicle stress testing. *Simulation Modelling Practice and Theory*.
49. Takahashi, S., et al. (2025). Cloud-edge synchronization protocols for real-time traffic digital twins. *IEEE Internet of Things Journal*.
50. O'Connor, L., & Murphy, P. (2024). Ethical validation of driving algorithms within virtual environments. *AI and Ethics*.
51. Qin, Y., & Zhou, B. (2026). Scaling digital twin simulations utilizing distributed post-quantum computing. *Future Internet*.
52. Al-Qahtani, F., et al. (2025). Reinforcement learning from human feedback (RLHF) in simulated driving environments. *Machine Learning*.
53. Gomez-Perez, A., & Ruiz, C. (2024). Benchmarking edge-AI algorithms for real-time vehicular control. *Journal of Systems Architecture*.
54. Nakamura, T., & Sato, Y. (2026). Quantitative analysis of QoS in 6G semantic communications. *IEEE Communications Letters*.



55. Li, W., & Fang, Z. (2025). Metrics for explainability in neuro-symbolic medical AI. *Artificial Intelligence in Medicine*.
56. Gupta, S., et al. (2024). High-fidelity federated learning for decentralized MRI analysis. *Medical Image Analysis*.
57. Robinson, D., & Patel, K. (2026). Privacy-preserving multi-institutional collaborative AI for oncology. *Nature Medicine*.
58. Chen, T., & Wang, S. (2025). Power minimization in RIS-assisted 6G networks via alternating optimization. *IEEE Transactions on Communications*.
59. Al-Dhahir, N., et al. (2024). Holographic MIMO and RIS: A physical layer perspective for 6G. *IEEE Wireless Communications*.
60. Zhao, L., & Zheng, M. (2026). A survey on the carbon footprint of training large language models for network optimization. *Sustainable Computing: Informatics and Systems*.