

Living Federated Digital Twins: Self-Evolving Digital Twins for Autonomous Wireless Networks

Mohammad Al-Quraan*, Ahsan R. Khan*, Ahmed Zoha*, Ahmad Taha*, Stephen McQuistin[†],
Buse Bilgin[‡], Adrian Kliks[§], and Paul Harvey[¶]

^{*}James Watt School of Engineering, University of Glasgow, Glasgow, UK

[†]School of Computer Science, University of St Andrews, St Andrews, UK

[‡]Turkcell, Istanbul, Turkey

[§]Poznan University of Technology and Rimedo Labs, Poznan, Poland

[¶]School of Computing Science, University of Glasgow, Glasgow, UK

Emails: {Mohammad.Alquraan, Ahsan.Khan, Ahmed.Zoha, Ahmad.Taha}@glasgow.ac.uk,
sjm55@st-andrews.ac.uk, buse.bilgin@turkcell.com.tr,
adrian.kliks@put.poznan.pl, Paul.Harvey@glasgow.ac.uk

Abstract—Next-generation wireless networks are expected to operate in highly dynamic environments characterized by dense deployments, distributed intelligence, and continuously changing radio conditions. Digital twins (DTs) have emerged as a promising tool for network monitoring, optimization, and automation. However, most existing DT implementations remain static and rely on centralized data collection and model training. Such designs limit their ability to adapt to evolving network conditions while raising concerns regarding scalability, communication overhead, and data privacy. In this paper, we introduce Living Federated Digital Twins (LFDTs), an architecture that allows DTs to continuously evolve through distributed learning. The proposed framework is built around DT objects (DTOs), modular representations of network entities that collaboratively learn through federated learning (FL) without exchanging raw data. We demonstrate the feasibility of the concept through a beam and blockage prediction case study using multimodal wireless and visual data. Results show that continuous updates enable beam prediction accuracy to improve from 60.4% to 91.0% and blockage prediction accuracy from 59.9% to 97.5%, while achieving performance close to centralized learning. These findings show the potential of LFDTs as a foundation for adaptive and intelligent network management in future wireless systems.

Index Terms—Digital twins, federated learning, next-generation wireless networks, blockage prediction.

I. INTRODUCTION

The vision of future wireless networks is increasingly centered on autonomy. Rather than being configured and optimized manually, network functions are expected to monitor their operating environment, anticipate changes, and adapt in real time. Achieving this requires accurate representations of network behavior that can support analysis and decision making under continuously changing conditions.

Digital twins (DTs) have emerged as a promising mechanism for enabling such capabilities. By maintaining a virtual representation of a physical network, a DT can support performance monitoring, simulation, prediction, and optimization. Consequently, DTs are becoming an important

component of the broader vision for AI-native network operation [1].

However, a fundamental limitation exists in how DTs are commonly built today. Most DT implementations rely on periodically collected data and centrally trained models [2]. While such approaches can capture network behavior at a particular point in time, they struggle to remain synchronized with environments that evolve continuously. User mobility, environmental changes, traffic fluctuations, and infrastructure updates can rapidly reduce the accuracy of a static digital representation. In practice, the twin gradually becomes less representative of the system it is intended to mirror.

We argue that future network DTs should not be viewed as static digital replicas but as continuously evolving entities capable of learning throughout their operational lifetime. This shift requires two capabilities. First, learning must occur near the data source to avoid excessive data movement and privacy concerns. Second, the digital representation itself must evolve as new observations become available.

To address these requirements, this paper introduces living federated digital twins (LFDTs). LFDTs extend the traditional DT concept by combining continuous learning with federated intelligence. At the core of the framework are digital twin objects (DTOs), modular representations of network entities that maintain local knowledge while contributing to a shared network-wide DT. Through federated learning (FL), DTOs continuously refine their models using local observations without exposing raw data [3].

To demonstrate how such a framework can be realized, we present a beam and blockage prediction case study using distributed multimodal observations. The objective is not to propose a new prediction model, but rather to demonstrate how continuously evolving DTOs can collectively create and maintain a living DT of a wireless system.

The remainder of this paper is organised as follows. Section II discusses the limitations of conventional DTs and motivates the need for living DTs. Section III presents




Figure 1: LFDt workflow. Multi-modal observations are transformed into DTOs, which continuously evolve through FL and collectively support intelligent network management applications.

the proposed LFDt framework and introduces the concept of DTOs. Section IV demonstrates the applicability of LFDt through a beam and blockage prediction case study, while Section V outlines key research directions and open challenges. Finally, Section VI concludes the paper.

II. FROM STATIC DIGITAL TWINS TO LIVING DIGITAL TWINS

Traditional DTs have largely inherited the assumptions of conventional simulation environments [4]. A model is created, calibrated using available data, and then used to analyze system behavior. While effective in relatively stable environments, this paradigm becomes increasingly problematic where conditions continuously evolve. Recent advances in 6G, AI-native networking, and autonomous network management have significantly increased the expectations placed on digital representations of wireless systems [5], [6]. Modern networks are no longer required only to support performance monitoring and offline analysis, but are increasingly expected to enable proactive optimization, automated decision-making, and closed-loop control. To fulfil these requirements, DTs must maintain an accurate understanding of both the current network state and its likely future evolution. However, achieving this level of awareness becomes increasingly challenging as wireless environments become more dynamic and heterogeneous.

Existing DT implementations often rely on periodic synchronization between physical and digital entities. While this approach is sufficient for relatively static systems, it introduces a growing mismatch between the digital representation and the physical network when conditions change rapidly. In dense wireless deployments, network conditions may evolve within seconds due to user mobility, blockage

events, traffic fluctuations, and radio resource reconfiguration. Consequently, a DT constructed using historical observations may progressively lose its ability to accurately represent the operating environment, reducing its value for prediction and decision support [1], [7]. The challenge becomes even more pronounced when considering the scale of future wireless networks. Emerging deployments are expected to incorporate distributed intelligence, edge computing resources, sensing capabilities, and large collections of interconnected devices [8], [9]. These systems continuously generate large volumes of heterogeneous data that must be processed and incorporated into the digital representation in a timely manner.

Traditional centralized workflows struggle to accommodate such requirements due to communication costs, privacy considerations, and the complexity of integrating multi-modal information generated across geographically dispersed locations. Three challenges motivate a different approach. First, wireless environments are inherently dynamic. User mobility, environmental blockages, changing traffic patterns, and infrastructure updates continuously alter network behavior. As a result, a DT built from historical data can quickly become outdated. Second, modern wireless systems generate data across geographically distributed entities. Collecting all observations at a centralized location introduces communication overhead and creates potential privacy and ownership concerns [10]. Third, future networks increasingly rely on heterogeneous data sources, including radio measurements, network telemetry, positioning information, cameras, and environmental sensors. Integrating these diverse observations into a single centralized workflow becomes increasingly difficult as network scale grows.

These limitations suggest the need for a new DT paradigm that is not only data-driven but also continuously adaptive.




Figure 2: Modular LFDT architecture based on federated DTOs. Each DTO maintains a localized representation of network behavior while contributing to a larger network-wide DT.

Rather than viewing a DT as a periodically updated digital replica, future wireless systems require a living digital representation that evolves together with the physical environment, continuously incorporates new knowledge, and supports collaborative learning across distributed network entities. Instead of periodically reconstructing a DT, the representation should evolve incrementally as new observations arrive. Furthermore, the learning process should remain distributed, allowing network entities to collaborate without exchanging raw data. We refer to this paradigm as an LFDT.

III. LIVING FEDERATED DIGITAL TWINS

The proposed LFDT framework is illustrated in Fig. 1. LFDT builds upon recent advances in network DTs and federated intelligence to support continuously evolving network representations [11], [12]. LFDT is designed around a simple principle: the digital representation must evolve together with the physical system it represents. Instead of maintaining a single monolithic twin, LFDT decomposes the network into a collection of collaborative digital entities. Each entity continuously learns from local observations while simultaneously benefiting from knowledge gathered elsewhere in the network. This design enables scalability, adaptability, and privacy preservation while maintaining a coherent network-wide representation.

A. Digital Twin Objects

The modular structure of the proposed LFDT architecture is illustrated in Fig. 2. Rather than representing an entire wireless network as a single monolithic DT, LFDT decomposes the system into a collection of DTOs [13]. Each DTO maintains a localized representation of a network entity, function, or geographical region while remaining connected to the broader DT ecosystem. A network-wide LFDT therefore emerges from the federation and composition of multiple DTOs rather than from a single centralized model.

DTOs are the fundamental building blocks of the LFDT framework. A DTO may represent an individual base station,

a cluster of cells, a mobility management function, or another network subsystem depending on the application requirements. This flexibility allows network operators to define digital representations at different levels of granularity while maintaining a common framework for collaboration and orchestration. Each DTO captures two complementary aspects of the physical system:

- Structural information, including topology, coverage characteristics, service configuration, and resource allocation policies.
- Behavioral information, including user mobility, beam selection dynamics, interference patterns, traffic evolution, and other operational characteristics.

The key distinction between a DTO and a conventional DT component is continuous evolution. As new observations become available, DTOs update their internal models and refine their understanding of the local environment. This allows the digital representation to remain aligned with changing network conditions rather than becoming a static snapshot of historical behavior.

Multiple DTOs can then be composed to form a larger network-wide LFDT. Through this compositional approach, the framework remains scalable while allowing different parts of the network to evolve independently and at different rates.

B. Continuous Federated Evolution

The defining characteristic of LFDT is its ability to continuously evolve. Each DTO collects observations from its local environment and updates its internal learning model. Rather than sharing raw data, DTOs exchange model updates through an FL process [14]. An aggregation entity combines these updates into an improved shared model that is redistributed to participating DTOs. This process creates a closed learning loop in which both local DTOs and the global DT continuously improve over time. Consequently, the twin becomes a living representation of the physical network rather than a static snapshot generated from historical observations.

IV. CASE STUDY: BEAM AND BLOCKAGE PREDICTION

To demonstrate the feasibility of LFDT, we consider beam and blockage prediction in a mmWave wireless environment. The study uses the ViWi framework [15], [16], which provides synchronized visual and wireless observations generated within a simulated urban environment. Beam and blockage prediction have emerged as important enablers of reliable high-frequency wireless communications [17]. Three base stations are represented as independent DTOs, each maintaining local observations and participating in FL training. A total of 5,000 samples are distributed across the DTOs, while an independent test set is used for evaluation. To emulate the evolution of a real deployment, data are introduced through five consecutive temporal intervals. At each interval, DTOs update their local models and participate in federated aggregation. This experimental setup allows us to investigate a central question: *can a DT improve as the underlying physical*

Table I: Comparison of centralized learning, isolated local learning, and LFDT-based FL.

Method	Blockage (%)	Beam Top-1 (%)
Centralised	98.22	93.69
Local (per DTO)	94.44	82.81
Federated (LFDT)	97.50	91.02

environment evolves? The answer provides insight into the practical feasibility of living DTs.

A. Learning Model and Federated Configuration


Each DTO employs a lightweight convolutional neural network (CNN) to jointly process visual and wireless observations. The model consists of a two-dimensional convolutional layer with 16 filters and a 3×3 kernel, followed by a max-pooling layer, a flattening stage, and a fully connected layer with 100 neurons using the rectified linear unit (ReLU) activation function. A dropout layer is incorporated to mitigate overfitting. The chosen architecture intentionally remains lightweight to reflect the computational constraints typically encountered at edge-deployed DTOs. The output layer uses a softmax activation for beam prediction and a sigmoid activation for blockage prediction.

The model is trained locally using the Adam optimizer with a learning rate of 0.001 and a batch size of 32. During each FL round, every DTO performs five local training epochs before transmitting model parameters to the federation server. Global aggregation is performed using the Federated Averaging (FedAvg) algorithm, where the contribution of each DTO is weighted according to its local dataset size [18]. A total of 40 communication rounds are performed during each temporal interval.


To emulate realistic network heterogeneity, the dataset is partitioned across the three DTOs according to user trajectories, creating a non-independent and non-identically distributed (non-IID) learning environment. The temporal evolution of the LFDT is modelled through five consecutive training intervals. At each interval, newly available observations are incorporated into local DTO training before participating in federated aggregation. This process allows the LFDT to continuously refine its representation of the environment and capture evolving network conditions over time.

B. Results and Discussion

To assess the effectiveness of LFDT, we compare its performance against two baseline configurations. In the *centralized* setup, training data from all DTOs are aggregated at a single location and used to train a global model. In the *local* setup, each DTO trains independently using only its local observations without exchanging information with other DTOs. The reported local accuracy corresponds to the average performance across the three DTOs. These baselines provide insight into the trade-off between collaborative learning and data locality.



(a)



(b)

Figure 3: LFDT prediction accuracy across five temporal updates: (a) beam prediction and (b) blockage prediction.

Table I summarizes the performance of centralized learning, isolated local training, and the proposed LFDT framework. LFDT achieves performance close to centralized learning for both prediction tasks while preserving data locality and avoiding raw data exchange. For blockage prediction, LFDT reaches 97.5% accuracy compared with 98.2% under centralized training. Similarly, LFDT achieves 91.0% beam prediction accuracy, remaining within 2.7 percentage points of the centralized baseline. These results indicate that federated DTO collaboration can retain most of the predictive capability of centralized training while offering a more scalable and privacy-preserving solution. These results validate the effectiveness of the federated architecture. However, they do not fully illustrate the defining characteristic of LFDT, namely its ability to evolve over time. To investigate this behaviour, Fig. 3 presents the prediction accuracy obtained across five consecutive temporal updates.

As shown in Fig. 3, the LFDT begins with a limited representation of the environment during the first training interval, resulting in beam prediction and blockage prediction accuracies of 60.4% and 59.9% respectively. As additional observations become available and successive federated updates are performed, both tasks exhibit consistent performance

improvements. By the fifth temporal update, beam prediction accuracy reaches 91.0%, while blockage prediction accuracy increases to 97.5%. This trend highlights the key distinction between LFDT and conventional DT implementations. Rather than relying on a static model trained once and periodically reconstructed, LFDT continuously refines its internal representation using new observations collected from distributed DTOs. The progressive improvements shown in Fig. 3 demonstrate that the twin becomes increasingly representative of the physical environment as operational knowledge accumulates.

Taken together, the results in Table I and Fig. 3 suggest that LFDT provides two complementary benefits. First, it achieves predictive performance comparable to centralized learning without requiring raw data sharing. Second, and more importantly, it enables a continuously evolving digital representation capable of adapting to changes in network conditions over time. This adaptive behaviour is central to the concept of a living DT and forms the foundation for future autonomous network management systems.

V. FUTURE RESEARCH DIRECTIONS

Although the presented case study demonstrates the potential of LFDTs, several research challenges must be addressed before they can be deployed at scale in operational wireless networks. One important research direction concerns *automated DTO creation and composition*. In this work, DTOs are manually defined based on network entities and deployment requirements. Future LFDT implementations should be capable of automatically identifying suitable network partitions, instantiating DTOs, and adapting their granularity as network conditions evolve [19]. Such capabilities would enable scalable and self-organising DT ecosystems capable of supporting increasingly complex wireless infrastructures.

Another challenge relates to *learning from heterogeneous and non-IID data*. Wireless environments naturally produce diverse data distributions due to differences in mobility patterns, traffic demands, deployment characteristics, and environmental conditions. While FL enables collaboration across distributed DTOs, significant statistical heterogeneity may affect convergence and model performance [10]. Future research should investigate personalised FL, transfer learning, and adaptive aggregation mechanisms to improve robustness under highly diverse network conditions.

The *communication efficiency* of LFDTs also requires further investigation. Although FL significantly reduces the need for raw data exchange, model synchronisation may still introduce substantial overhead as the number of DTOs increases [20]. Future work should explore model compression, sparse updates, asynchronous federation, and hierarchical aggregation strategies to reduce communication costs while preserving model accuracy.

Finally, a fundamental challenge lies in enabling *continual and adaptive learning* throughout the operational lifetime of an LFDT. As network conditions evolve, DTOs must continuously incorporate new observations while retaining

previously acquired knowledge. Developing mechanisms that balance adaptability with long-term model stability, while avoiding issues such as model drift and catastrophic forgetting, will be critical to fully realising the vision of living DTs. Addressing these challenges will be essential for practical LFDT deployment in future 6G networks.

VI. CONCLUSIONS

This paper introduced living federated digital twins (LFDTs), a framework for creating continuously evolving digital representations of wireless networks. Unlike conventional DTs that rely on static models and centralized data collection, LFDT combines modular DTOs with FL to enable distributed, privacy-preserving, and adaptive digital twinning. Through a beam and blockage prediction case study, we demonstrated how multiple DTOs can collaboratively construct and refine a network-wide DT while continuously learning from new observations. The results showed that LFDT achieves performance close to centralized learning and improves prediction accuracy over time. These findings highlight the potential of LFDTs to support intelligent and autonomous network management in future wireless systems. Future work will focus on automated DTO orchestration, communication-efficient federation, and lifelong learning mechanisms that further enhance the adaptability and scalability of living DTs.

REFERENCES

- [1] L. U. Khan, Z. Han, W. Saad, E. Hossain, M. Guizani, and C. S. Hong, "Digital twin of wireless systems: Overview, taxonomy, challenges, and opportunities," *IEEE Communications Surveys & Tutorials*, 2022.
- [2] A. Alkhateeb, S. Jiang, and G. Charan, "Real-time digital twins: Vision and research directions for 6G and beyond," *IEEE Communications Magazine*, vol. 61, no. 11, pp. 128–134, 2023.
- [3] M. Al-Quraan *et al.*, "Edge-native intelligence for 6G communications driven by federated learning: A survey of trends and challenges," *IEEE Transactions on Emerging Topics in Computational Intelligence*, vol. 7, no. 3, pp. 957–979, Mar. 2023.
- [4] C. Semeraro, M. Lezoche, H. Panetto, and M. Dassisi, "Digital twin paradigm: A systematic literature review," *Computers in Industry*, vol. 130, p. 103469, 2021.
- [5] J. Hoydis, F. A. Aoudia, A. Valcarce, and H. Viswanathan, "Toward a 6G AI-Native air interface," *IEEE Communications Magazine*, vol. 59, no. 5, pp. 76–81, 2021.
- [6] P. Harvey, A. Tatar, P. Imai, L. Wong, and L. Bringuier, "Evolutionary autonomous networks," *Journal of ICT Standardization*, pp. 201–228, 2021.
- [7] P. Almasan, M. Ferriol-Galmés, J. Paillisse, J. Suárez-Varela, D. Perino, D. López, A. A. P. Perales, P. Harvey, L. Ciavaglia, L. Wong *et al.*, "Network digital twin: Context, enabling technologies, and opportunities," *IEEE Communications Magazine*, vol. 60, no. 11, pp. 22–27, 2022.
- [8] D. C. Nguyen, M. Ding, P. N. Pathirana, A. Seneviratne, J. Li, D. Niyato, O. Dobre, and H. V. Poor, "6G internet of things: A comprehensive survey," *IEEE Internet of Things Journal*, vol. 9, no. 1, pp. 359–383, 2021.
- [9] H. Zhang, N. Shlezinger, F. Guidi, D. Dardari, M. F. Imani, and Y. C. Eldar, "Near-field wireless power transfer for 6G internet of everything mobile networks: Opportunities and challenges," *IEEE Communications Magazine*, vol. 60, no. 3, pp. 12–18, 2022.
- [10] Q. Li, Z. Wen, Z. Wu, S. Hu, N. Wang, Y. Li, X. Liu, and B. He, "A survey on federated learning systems: vision, hype and reality for data privacy and protection," *IEEE Transactions on Knowledge and Data Engineering*, 2021.
- [11] L. U. Khan, W. Saad, D. Niyato, Z. Han, and C. S. Hong, "Digital-twin-enabled 6G: Vision, architectural trends, and future directions," *IEEE Communications Magazine*, vol. 60, no. 1, pp. 74–80, 2022.

- [12] S. P. Ramu, P. Boopalan, Q.-V. Pham, P. K. R. Maddikunta, T. Huynh-The, M. Alazab, T. T. Nguyen, and T. R. Gadekallu, "Federated learning enabled digital twins for smart cities: Concepts, recent advances, and future directions," *Sustainable Cities and Society*, vol. 79, p. 103663, 2022.
- [13] "Autonomous Networks -Architecture Framework," United Nations International Telecommunication Union, Geneva, CH, Standard, Dec. 2023.
- [14] Q. Yang, Y. Liu, T. Chen, and Y. Tong, "Federated machine learning: Concept and applications," *ACM Trans. Intell. Syst. Technol.*, vol. 10, no. 2, Jan. 2019. [Online]. Available: <https://doi.org/10.1145/3298981>
- [15] M. Alrabeiah, A. Hredzak, Z. Liu, and A. Alkhateeb, "Viwi: A deep learning dataset framework for vision-aided wireless communications," in *2020 IEEE 91st Vehicular Technology Conference (VTC2020-Spring)*, 2020, pp. 1–5.
- [16] M. Al-Quraan *et al.*, "A hybrid data manipulation approach for energy and latency-efficient vision-aided UDNs," in *Eighth International Conference on Software Defined Systems (SDS)*, Gandia, Spain. IEEE, 2021, pp. 1–7.
- [17] Al-Quraan, Mohammad *et al.*, "Federated learning for reliable mmwave systems: Vision-aided dynamic blockages prediction," in *IEEE Wireless Communications and Networking Conference (WCNC)*, Glasgow, UK. IEEE, 2023, pp. 1–6.
- [18] M. Al-Quraan, A. Khan, A. Centeno, A. Zoha, M. A. Imran, and L. Mohjazi, "Fedratrees: A novel computation-communication efficient federated learning framework investigated in smart grids," *Engineering Applications of Artificial Intelligence*, vol. 124, p. 106654, 2023.
- [19] S. Ding, D. Flynn, and P. Harvey, "Automatic generation of digital twins for network testing," in *2025 IEEE 45th International Conference on Distributed Computing Systems Workshops (ICDCSW)*, 2025.
- [20] D. Wu, R. Ullah, P. Harvey, P. Kilpatrick, I. Spence, and B. Varghese, "Fedadapt: Adaptive offloading for IoT devices in federated learning," *IEEE Internet of Things Journal*, vol. 9, no. 21, pp. 20 889–20 901, 2022.