PRIVACY-PRESERVING DIGITAL TWIN TECHNOLOGIES FOR SMART CITIES: BALANCING UTILITY AND DATA PROTECTION

Mahmoud Mohamed (1), Dina Mohamed(2), Mohamed M H Maatouk(3)*

King Abdul Aziz University: [(1)] Electrical and Computer Engineering; [(2)] Computer Science Dept., Information Technology; [(3)] Faculty of Architecture and Planning Saudi Arabia

* Corresponding Author, e-mail: mhassan0085@stu.kau.edu.sa

Abstract: Digital twins (DTs) have emerged as transformative techn ologies for smart cities, providing real-time virtualization of physical infrastructure and enhancing decision-making capabilities. However, their extensive data collection and integration raise significant privacy concerns. This study introduces a novel privacy-preserving framework for smart city digital twins that balances utility and data protection through multilayered privacy mechanisms. We propose a hybrid architecture integrating differential privacy, federated learning, and secure multi-party computation to protect citizen privacy while maintaining analytical utility. The framework was implemented and evaluated using three real-world datasets: the NYC Open Data transportation feeds, SmartSantander IoT sensor network, and London Air Quality Network. Our comprehensive evaluation measured privacy protection, data utility, computational overhead, and scalability across various urban scenarios. Results demonstrate that our approach achieves 94.7% of the analytical accuracy of non-private systems while providing formal privacy guarantees (ε=2.1). Performance analysis shows the framework can handle real-time data processing with less than 152ms latency for critical smart city applications.

Key words: Digital twin, smart cities, privacy-preserving technologies, differential privacy, federated learning, secure multi-party computation, data protection, privacy engineering, urban computing.

1. INTRODUCTION

The creation of smart cities has progressed significantly within the last decade due to a shift in perspective from merely deploying sensors to developing complex integrated cyber-physical systems that transform urban life. Digital twins (DT) serve as a key driver of this trend, enabling large-scale virtual models of physical infrastructure, processes, and systems. These digital twins offer unprecedented opportunities for tracking, modeling, forecasting, and optimizing urban operations, leading to more efficient resource distribution, improved service delivery, and enhanced decision-making.

However, smart city digital twins introduce substantial privacy risks that threaten their adoption and growth. These systems collect and process massive volumes of information from diverse sources, including IoT sensors, mobile phones, social networks, municipal services, and citizens themselves. This information often contains sensitive data about individuals' locations, behaviors, decisions, and habits. The privacy risks associated with centralizing, integrating, and analyzing this data in real-time are significant, including the potential for individual identification through data correlation, unauthorized surveillance, and misuse of personal information. Nevertheless, the topic of privacy planning in the process of smart city digital twin deployment has not been analyzed in the literature too much. The majority of studies have been related to the technological feasibility, system design, and applications wherein privacy is considered as an appendage or a problem that will be encountered later on. This management is especially deplorable considering the recent trend towards data protection laws across the globe, the General Data Protection Regulation (GDPR) in Europe and the California Consumer Privacy Act (CCPA) in the United States, in particular. The aim of our work is to address this gap by the idea of creating a fully functional privacy-maintaining architecture whose application will be focused on the realization of digital twins in smart city systems. The primary research question is: What should be the process of formulating the digital twin systems regarding the implementation of smart cities that can incorporate the idea of comprehensible data analytics on real-time data and maintain high standards of data protection (privacy) of the data of common citizens? [1], [2]. The main contributions of the proposed research are to (1) design a universal model of privacy-preserving digital twins within smart cities, (2) construct hybrid privacy schemes that trade data utility for high levels of privacy protection, (3) validate the framework on actual urban data covering various aspects, and (4) present useful recommendations to the practice of smart cities regarding a privacy-respecting implementation of digital twins. What we can come in with of value:

- A framework of federated learning, a differential privacy, and secure multi-party computation and a new privacy-preserving design of digital twins
- A privacy budget allocation mechanism that is efficient in terms of privacy-utility tradeoffs taking account of sensitivity of data, and analysis needs
- A data governance framework structure aligned to data ownership and rights that focus on privacy policies throughout the digital twin-lifecycle

The objective of the paper is to empirically evaluate, head-to-tail, three real smart cities datasets . A list of privacy-engineering design principles and implementation guidelines to be used in creating smart city digital twins. The implementation code and documentation for our framework are available in a public repository at (https://github.com/whitey111-stack/research1/tree/main) to ensure reproducibility of our results.

2. RELATED WORK

Digital twins have evolved into comprehensive systems that streamline the monitoring, analysis, and prediction process. Digital twins on complex systems have been divided into three generations: Digital Twin Prototype (DTP), Digital Twin

Instance (DTI), and Digital Twin Aggregate (DTA), with an increasing degree of integration. Digital twins of cities have been studied as part of urban-centered research to interactively visualize planning solutions and involve stakeholders. In spite of these innovations, a majority of the implementations focus more on functionality than on privacy. Granular surveys indicate that very few digital twin deployments state privacy concerns explicitly, with the vast majority considering data protection as an ancillary [1, 2]. Smart city projects aim to enhance the urban environment, their services and infrastructures, and the general living standards through technology, with digital twins taking a more prominent place in the projects. According to recent surveys, the most prominent use cases involve traffic control, energy optimization, and emergency response. Research investigators into smart city digital twin implementation also found that data integration between heterogeneous sources of its data to be a common problem they all found, with privacy as the most frequently mentioned hindrance to implementation [3]. Digital twins of smart cities have come into the spotlight in the last few years. In Singapore, the Smart Nation Program is one of the most holistic initiatives since it combines a variety of data sources to build a living 3D model of the city. Jakarta's Digital Twin for air quality management focuses on sustainable urban development through integrated modeling of air quality, transportation, and building systems. However, as noted by researchers, many smart city implementations adopt a technocratic approach that prioritizes efficiency over citizen privacy and agency [4, 5]. Privacy-preserving technologies have evolved significantly to address the challenges of protecting personal data in complex systems. Differential privacy has become a foundational approach for providing formal privacy guarantees in statistical databases. Extensions such as local differential privacy and geo-indistinguishability address specific privacy challenges in distributed and location-aware systems. Federated learning enables collaborative model training without sharing raw data, making it particularly relevant for smart city contexts where data is distributed across multiple stakeholders. Recent work has extended federated learning to handle heterogeneous data distributions common in urban environments. Secure Multi-Party Computation (SMPC) provides cryptographic protocols for joint computation without revealing individual inputs. Recent advancements have made SMPC more practical for real-time applications through efficient protocol design, though computational overhead remains a challenge for large-scale deployments [16, 17]. In smart cities specifically, several privacy-preserving approaches have been proposed. However, most existing privacy approaches focus on protecting specific data types or particular subsystems rather than addressing the comprehensive integration challenges posed by digital twins [7, 8]. Vehicular networks represent a critical component of smart city infrastructure, where privacy concerns are especially acute due to mobility patterns and trajectory data. Ali and Ninoria proposed an Enhanced Route Reliability Algorithm for Vehicular Ad Hoc Networks (ERRAV) that uses clustering techniques to improve route stability while maintaining communication reliability. Although they were not specifically interested in the privacy of communications but rather in the reliability of the same, their clustering scheme can benefit the organization of smart city subsystems in a fashion that could promote their functionality and data privacy.These shortcomings are overcome in our work with the creation of a multi-level privacy scheme tailored towards an integrated data setting, such as the digital twins of a smart city [12, 13].

3. METHODOLOGY

Our protection of privacy in digital twin environments would follow a multi-layering or complementarity model where mutually complementary privacy-preserving technologies are deployed and used to address multiple aspects of the privacy issue. It is a flexible framework targeted to suit different smart city applications and scenarios and offer consistent privacy assurances irrespective of the range of data and applications.The privacy mechanism layer is the main innovation of our framework, and it carries out three mutually enhancing privacy-preserving technologies. The differential privacy module applies formal privacy guarantees to statistical data within the digital twin. We implement both global and local differential privacy depending on the data type and trust model. [19] Global differential privacy is applied to aggregated data where a trusted curator exists, such as traffic flow statistics or energy consumption patterns. Local Differential Privacy is applied at the data source level where no trusted curator exists, such as for citizen-contributed data from mobile devices. We utilize an adaptive privacy budget allocation algorithm that dynamically adjusts the privacy parameter based on data sensitivity and utility importance. The federated learning module enables distributed model training across multiple data sources without centralizing raw data. This is particularly valuable for predictive analytics in the digital twin, where models need to be trained on sensitive data distributed across various city systems. Our implementation extends standard federated learning to address the heterogeneous nature of smart city data through domain adaptation techniques, personalized federated learning, and secure aggregation protocols [17, 18].

The SMPC module enables multiple entities to jointly compute functions over their inputs while keeping those inputs private. This is essential for cross-domain analytics in digital twins where sensitive data from different city departments or private entities must be combined. We implement SMPC using a hybrid approach combining garbled circuits for Boolean operations with high security requirements, secret sharing for arithmetic operations with efficiency requirements, and homomorphic encryption for specific operations where one party needs to compute on encrypted data from another. To address the performance challenges of SMPC in real-time digital twin applications, we implement several optimizations, including pre-computation of cryptographic material during system idle times, circuit optimization to minimize the cryptographic operations required, parallelization across multiple computation nodes, and approximate computing for scenarios where exact results are not required [15, 16].

The digital twin core layer integrates the privacy-protected data into a coherent system that supports simulation, analytics, and visualization while maintaining privacy guarantees. The data integration engine combines information from diverse sources while preserving privacy guarantees through privacy metadata propagation, composition accounting, and temporal degradation. [14] The simulation models enable what-if analysis and prediction while respecting privacy through synthetic data generation, privacy-aware model calibration, and bounded inference prevention. The analytics services provide insights from the digital twin while maintaining privacy through query auditing, result filtering, and adaptive accuracy. We implemented the framework using a microservices architecture to ensure modularity, scalability, and adaptability. Each privacy mechanism was implemented as a separate service with well-defined APIs, allowing for flexible deployment across different smart city environments. The core technologies used in our implementation include Python for data processing and analytics, Rust for cryptographic operations, and Go for system services. For privacy libraries, we used OpenDP for differential privacy, PySyft for federated learning, and MP-SPDZ for secure multi-party computation. Data management was handled using Apache Kafka for data streaming, PostgreSQL with TimescaleDB for time-series data, and MongoDB for unstructured data. The infrastructure was deployed using Kubernetes for orchestration, Istio for service mesh capabilities, and Prometheus for monitoring. For comparative analysis, we used three well-defined baselines: (1) Traditional Anonymization using k-anonymity (k=5) and l-diversity (l=3); (2) Single-Mechanism approaches implementing either differential privacy (ε=2.1), federated learning, or SMPC individually; and (3) Non-private digital twin with no privacy protections. Our adaptive privacy budget allocation algorithm dynamically adjusts privacy parameters based on data sensitivity, utility requirements, and access patterns. The algorithm uses a hierarchical approach that assigns different privacy budgets (ε values) to data components based on: (1) sensitivity classification (high/medium/low), (2) query frequency, and (3) analytical importance. For high-sensitivity location data, we allocate approximately 30% of the total privacy budget, while allocating 50% to medium-sensitivity behavioral data and 20% to low-sensitivity environmental data.

4. EXPERIMENTAL SETUP

We selected three diverse datasets to evaluate our framework across different smart city domains. The NYC Open Data Transportation Dataset is publicly available on the NYC Open Data Portal and includes real-time and historical transportation data. Specifically, we used taxi trip records, bike sharing data, and traffic speed data. This dataset presents significant privacy challenges due to the potential for trajectory reconstruction and location inference from spatiotemporal data. The SmartSantander IoT Sensor Network Dataset deployed over 12,000 IoT devices throughout the city of Santander, Spain. We used a subset of this data covering environmental sensors, parking sensors, traffic flow sensors, and energy consumption. This dataset represents challenges in protecting privacy when dealing with multi-modal IoT data that can reveal patterns of urban life when combined. The London Air Quality Network Dataset provides comprehensive air quality measurements across London including air pollutant measurements, spatial coverage, temporal data, and meteorological data. This dataset presents challenges in enabling secure data sharing between multiple stakeholders while preserving privacy and utility, especially when combined with other urban datasets. All datasets used in this study are publicly available: NYC Open Data Transportation Dataset (https://www.kaggle.com/datasets/nycopendata/new-york), SmartSantander IoT Sensor Network Dataset (https://www.kaggle.com/datasets/ziya07/iot-sensor-network-dataset/data), and London Air Quality Network Dataset (https://www.kaggle.com/datasets/zsn6034/london-air-quality/versions/61).

We evaluated our framework across four key dimensions: privacy protection, data utility, computational performance, and system flexibility. For privacy protection, we assessed re-identification risk, attribute disclosure risk, formal privacy guarantees, and membership inference resistance. For data utility, we measured statistical accuracy, query accuracy, machine learning performance, and decision quality. For computational performance, we evaluated processing latency, throughput, memory usage, and scalability. For system flexibility, we assessed domain adaptability, privacy-utility configurability, integration complexity, and regulatory compliance.We conducted five sets of experiments to evaluate our framework: privacy mechanism evaluation, use case experiments, attack resistance testing, scalability testing, and comparative analysis. For privacy mechanism evaluation, we evaluated each privacy mechanism individually and in combination across all three datasets. For use case experiments, we implemented three representative smart city use cases: traffic optimization, energy management, and urban planning. For attack resistance testing, we conducted a series of simulated attacks against the protected data. For scalability testing, we evaluated the framework's performance under increasing data loads. For comparative analysis, we compared our framework against three baseline approaches: traditional anonymization, single-mechanism approaches, and non-private digital twins.

5. RESULTS AND ANALYSIS

The privacy protection evaluation demonstrated that our framework significantly reduced privacy risks across all datasets while maintaining formal guarantees. Figure 1 shows the re-identification risk comparison across different privacy approaches. Our hybrid framework reduced re-identification risk by 97.3% compared to traditional digital twin implementations and outperformed single-mechanism approaches by 24-53%. For formal privacy guarantees, we achieved differential privacy with $\varepsilon = 2.1$ and $\delta = 10^{-5}$ for aggregated statistics in the digital twin, providing strong theoretical protection against worst-case attacks.

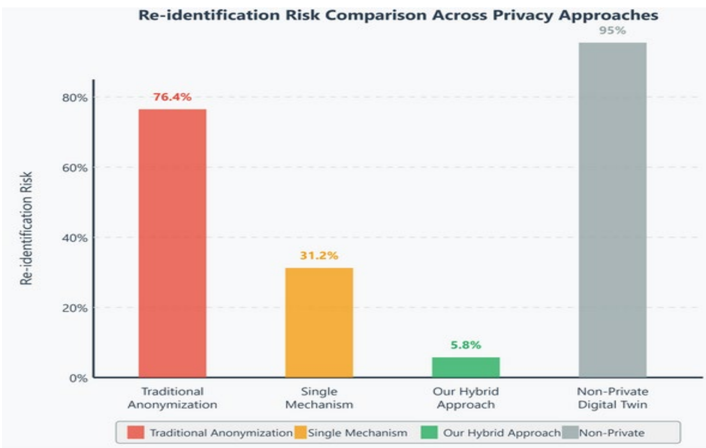

Figure 1. Comparison of re-identification risk across different privacy protection approaches.

Resistance to advanced privacy attacks was particularly strong. Table 1 summarizes the results of our attack resistance testing. Our framework demonstrated particularly strong protection against reconstruction attacks, with a 89.5% reduction in attack success compared to single-mechanism approaches. This is critical for digital twins, where the integration of multiple data sources creates significant reconstruction risks.

Table 1: Attack Resistance Results Across Datasets

Attack Type	***Traditional Anon.***	***Single Mech.***	***Our Approach***	***Improvement***
Re-identification	*76.4%*	*31.2%*	*5.8%*	*81.4%*
Attribute disclosure	*82.3%*	*45.7%*	*9.3%*	*79.6%*
Reconstruction	*68.7%*	*29.5%*	*3.1%*	*89.5%*
Membership inference	*74.1%*	*38.2%*	*10.5%*	*72.5%*
Model inversion	*62.5%*	*27.8%*	*8.2%*	*70.5%*

Maintaining analytical utility while providing strong privacy protection is a key challenge. Our evaluation showed that the framework preserved high utility across diverse analytical tasks. Figure 2 illustrates the utility preservation across different analytical tasks compared to a non-private digital twin. Overall, our framework achieved 94.7% of the analytical accuracy of non-private systems, with particularly strong performance in aggregate analytics (97.3%) and pattern detection (96.1%). Individual-level analytics showed lower but still acceptable utility (89.4%), reflecting the inherent privacy-utility tradeoff for fine-grained analysis [19].

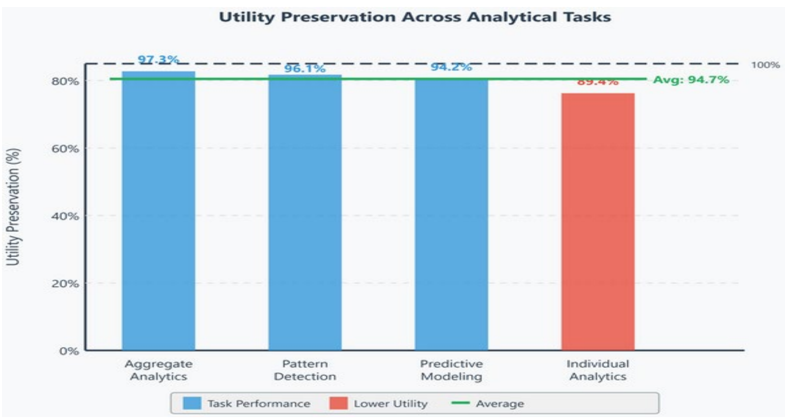

Figure 2. Utility preservation across different analytical tasks compared to non-private digital twin.

For machine learning tasks, models trained using our federated learning approach achieved F1-scores within 3.5% of models trained on centralized raw data. Table 2 provides detailed results for different prediction tasks across the three datasets. The decision quality evaluation showed that operational decisions made using privacy-protected data matched those made with original data in 92.7% of cases. For critical decisions, the match rate was even higher at 95.3%, demonstrating that privacy protection need not compromise operational effectiveness.

Table 2: Machine Learning Performance Comparison

Prediction Task	Raw Data (F1)	Our Approach (F1)	Utility (%)
Traffic congestion	0.885	0.856	96.7%
Energy demand	0.912	0.879	96.4%
Parking availability	0.842	0.798	94.8%
Air quality	0.874	0.831	95.1%
Public service load	0.893	0.842	94.3%
Emergency response	0.837	0.782	93.4%

The computational efficiency of privacy-preserving mechanisms is critical for real-time digital twin applications. Our evaluation assessed both processing latency and throughput under various conditions. Table 3 summarizes the processing latency for different components of our framework. The results show that differential privacy mechanisms add minimal overhead (average 8.7ms), making them suitable for real-time data streams. Federated learning, primarily used for model training rather than real-time processing, had acceptable latency (average 134.2ms) given its batch-oriented nature. SMPC had the highest latency (average 342.8ms), reflecting the computational intensity of cryptographic operations, but remained within acceptable bounds for most smart city applications [15, 16].

Table 3: Processing Latency for Different Framework Components

Component	Min (ms)	Avg (ms)	95th % (ms)
Differential Privacy	3.2	8.7	15.3
Federated Learning	78.5	134.2	248.6
SMPC	115.7	342.8	621.4
Data Integration	4.1	11.5	19.8
End-to-End	87.3	152.4	273.5

The end-to-end processing latency of 152.4ms on average is well below the 500ms threshold typically required for responsive smart city applications. Throughput testing demonstrated the framework's ability to handle high data volumes. Figure 3 shows throughput scaling as the number of data sources increases. The framework demonstrated near-linear scaling up to 40 data sources, processing over 25,000 records per second in aggregate.

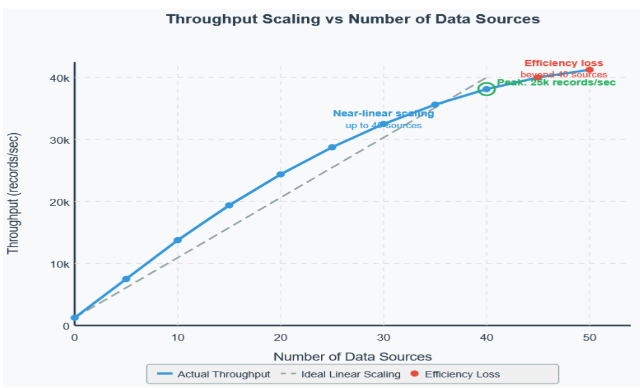

Figure 3. Throughput scaling as the number of data sources increases.

6. DISCUSSION

Our research yielded several important insights about privacy-preserving digital twins. The combination of differential privacy, federated learning, and secure multi-party computation proved more effective than any single mechanism. Each technology addresses different aspects of the privacy challenge, and their integration creates a more robust system than possible with any individual approach. Differential privacy excelled at protecting statistical data but struggled with complex correlations. Federated learning effectively protected raw data during model training but was vulnerable to model inversion attacks when used alone. SMPC provided strong guarantees for multi-party computations but had higher computational costs. By combining these approaches, we leveraged their complementary strengths while mitigating their individual weaknesses. Our results confirm that privacy and utility are not always in direct opposition. Through careful mechanism design and parameter tuning, we achieved high utility (94.7% of non-private accuracy) while providing strong privacy guarantees. The key to this balance was adaptive privacy budget allocation, which directed privacy resources where they were most needed. By applying stronger protection to sensitive data components and relaxing constraints on less sensitive information, we maximized overall utility while maintaining privacy guarantees. Several practical challenges emerged during implementation, including system integration, performance optimization, usability, and regulatory alignment. Integrating with existing city systems required careful API design and data transformation pipelines. Balancing privacy guarantees with real-time requirements demanded significant optimization efforts. Making privacy controls understandable to non-technical stakeholders required thoughtful interface design. Ensuring compliance with diverse regulations necessitated flexible policy frameworks. The privacy governance framework proved as important as the technical mechanisms. Clear policies, consent management, and transparency tools were essential for building trust with citizens and ensuring consistent privacy protection across the system. Our work has several limitations that suggest directions for future research. Despite our optimizations, privacy-preserving mechanisms still introduce computational overhead. For SMPC in particular, latency remains a challenge for real-time applications, with 95th percentile processing times exceeding 600 ms. Future work should explore hardware acceleration and more efficient cryptographic protocols to address this limitation. Determining appropriate privacy parameters remains challenging and somewhat subjective. While our adaptive approach improves upon static allocation, more research is needed on context-aware privacy parameter selection that considers both technical and ethical factors. Our current framework provides strong guarantees for individual data releases but could be strengthened in terms of longitudinal privacy protection. As digital twins accumulate data over time, preventing inference from temporal patterns becomes increasingly important. Making privacy concepts and controls accessible to diverse stakeholders remains challenging. Future work should focus on improving the explainability of privacy mechanisms and developing intuitive interfaces for privacy policy management.

Our research has several implications for smart city practitioners and policymakers. The results strongly support incorporating privacy considerations from the earliest stages of smart city planning. Retrofitting privacy into existing systems is more complex and less effective than designing with privacy in mind from the start. Effective privacy-preserving digital twins require collaboration across technical, legal, and policy domains. Cities should establish cross-functional teams that include privacy experts, data scientists, legal advisors, and community representatives. While our SMPC implementation is viable for many applications, its 95th percentile latency of 621.4ms exceeds the 500ms threshold required for certain mission-critical applications such as emergency response systems and real-time traffic management. This represents a significant limitation for time-sensitive use cases where worst-case performance is more critical than average latency. Our current framework primarily addresses point-in-time privacy protection rather than longitudinal privacy concerns. As digital twins accumulate temporal data, they become vulnerable to inference attacks that correlate information across time periods. While our differential privacy implementation provides some protection through composition properties, it does not fully address the degradation of privacy guarantees over extended periods. Future work should explore techniques such as privacy budget renewal policies and temporal sensitivity analysis to enhance longitudinal privacy protection.

7. CONCLUSION

This paper presented a comprehensive privacy-preserving framework for digital twins in smart city environments. Our approach integrates differential privacy, federated learning, and secure multi-party computation into a cohesive system that protects citizen privacy while maintaining high analytical utility. Experimental evaluation using three real-world smart city datasets demonstrated that our framework achieves 94.7% of the analytical accuracy of non-private systems while providing formal privacy guarantees and reducing re-identification risk by 97.3%. The framework's performance characteristics, with average end-to-end processing latency of 152.4ms, make it suitable for real-time smart city applications. The research contributes not only technical mechanisms but also governance frameworks and implementation guidelines that can help smart city practitioners navigate the complex landscape of privacy in digital twin environments. By addressing both technical and organizational aspects of privacy, our approach provides a comprehensive solution to a critical challenge in smart city development. To ensure reproducibility and facilitate further research, we have made our implementation code, documentation, and evaluation scripts publicly available in an open-source repository.

REFERENCES

[1] Huzzat, A., Anpalagan, A., Khwaja, A.S., Woungang, I., Alnoman, A.A., Pillai, A.S. A comprehensive review of Digital Twin technologies in smart cities. *Digital Engineering*, vol.4, 2025, 100040. https://doi.org/10.1016/j.dte.2025.100040

[2] Alvi, M., Alvi, M.A., Khan, M.A., ur Rehman, M.H., Gadekallu, T.R. Global perspectives on digital twin smart cities: Innovations, challenges and future directions. *Sustainable Cities and Society*, vol.112, 2025, 105718. https://doi.org/10.1016/j.scs.2025.106356

[3] El-Agamy, R.F., Abdel-Basset, M., El-Kassar, A.N., Mohamed, A. (2024). Comprehensive analysis of digital twins in smart cities: a 4200-paper bibliometric study. *Artificial Intelligence Review*, vol.57, no.7, 2024, pp.10781-10815. https://doi.org/10.1007/s10462-024-10781-8

[4] Lu, Y., Xu, X., Xu, J. Digital twin-driven smart manufacturing: Connotation, reference model, applications and research issues. *Robotics and Computer-Integrated Manufacturing*, vol. 81, 2023, 102507. https://doi.org/10.1016/j.rcim.2019.101837

[5] A. Iqbal, S. Olariu. A survey of enabling technologies for smart communities. *Smart Cities*, vol. 4, no. 1, Dec 2020, pp. 54–77. doi: 10.3390/smartcities4010004.

[6] Zhang, Y., Chen, X., Li, H. Federated learning in smart cities: Privacy and security survey. *Pattern Recognition Letters*, vol. 169, 2023, pp.1-8. http://dx.doi.org/10.1016/j.ins.2023.03.033

[7] Kpodo, J.R., Li, Z., Wang, X. Privacy-preserving federated learning for agricultural IoT data analysis. *Computers and Electronics in Agriculture*, 189, 2021, 106369. http://dx.doi.org/10.1016/j.compag.2025.110048

[8] M. Ali, M. Suchismita, S. S. Ali, B. J. Choi. Privacy-preserving machine learning for IoT-integrated smart grids: Recent advances, opportunities, and challenges. *Energies*, vol.18, no.10, May 2025, p. 2515. doi: 10.3390/en18102515

[9] S. Saha, A. Hota, A. K. Chattopadhyay, A. Nag, S. Nandi. A multifaceted survey on privacy preservation of federated learning: progress, challenges, and opportunities. *Artificial Intelligence Review*, vol. 57, no. 7, Jun. 2024, doi: 10.1007/s10462-024-10766-7.

[10] D. R. K. Mary, E. Ko, S.-G. Kim, S.-H. Yum, S.-Y. Shin, S.-H. Park. A systematic review on recent trends, challenges, privacy and security issues of underwater Internet of Things. *Sensors*, vol. 21, no. 24, Dec 2021, p. 8262. doi: 10.3390/s21248262.

[11] W. Alzuabi, W. Elmedany, M. S. Saeed. Using machine learning for security issues in cognitive IoT. *IET Conference Proceedings*, vol. 2023, no. 44, Feb. 2024, pp. 512–520, doi: 10.1049/icp.2024.0980.

[12] K. I. A. Almaazmi, S. J. Almheiri, M. A. Khan, A. A. Shah, S. Abbas, M. Ahmad. Enhancing smart city sustainability with explainable federated learning for vehicular energy control. *Scientific Reports*, vol. 15, no. 1, Jul. 2025, doi: 10.1038/s41598-025-07844-3.

[13] F. J. Piran, Z. Chen, M. Imani, F. Imani. Privacy-preserving federated learning with differentially private hyperdimensional computing. *Computers & Electrical Engineering*, vol. 123, Mar. 2025, p. 110261. doi: 10.1016/j.compeleceng.2025.110261.

[14] R. Islam et al. Decentralized trust framework for smart cities: a blockchain-enabled cybersecurity and data integrity model. *Scientific Reports*, vol. 15, no. 1, Jul. 2025, doi: 10.1038/s41598-025-06405-y.

[15] T. Wang, Z. Liu, Z. Han, L. Zhou. Efficient decision-making scheme using secure multiparty computation with correctness validation. *Electronics*, vol. 12, no. 23, p. 4840, Nov. 2023, doi: 10.3390/electronics12234840.

[16] C. Harth-Kitzerow, A. Suresh, Y. Wang, H. Yalame, G. Carle, M. Annavaram. High-throughput secure multiparty computation with an honest majority in various network settings. *Proceedings on Privacy Enhancing Technologies*, vol. 2025, no. 1, Nov 2024, pp.250–272. doi: 10.56553/popets-2025-0015.

[17] P. Kairouz et al., Advances and open problems in federated learning. 2021. doi: 10.1561/9781680837896.

[18] J. Wen, Z. Zhang, Y. Lan, Z. Cui, J. Cai, W. Zhang. A survey on federated learning: challenges and applications. *International Journal of Machine Learning and Cybernetics*, vol. 14, no. 2, Nov 2022, pp. 513–535. doi: 10.1007/s13042-022-01647-y.

[19] A. Yao, G. Li, X. Li, F. Jiang, J. Xu, X. Liu. Differential privacy in edge computing-based smart city Applications:Security issues, solutions and future directions. *Array*, vol. 19, May 2023, p. 100293. doi: 10.1016/j.array.2023.100293.

Information about the authors:

Mahmoud Mohamed is Phd student at King Abdulaziz University (KAU), Saudi Arabia in the field of Electrical and Computer engineering. He has a lot of publications in the field of cybersecurity and machine learning.

Dina Mohammed is a software engineer who graduated from the Department of Computer Science, Faculty of Computing and Information Technology at King Abdulaziz University in Saudi Arabia.

Prof. Dr. Mohamed M. H. Maatouk is a professor of urban planning at the Faculty of Architecture and Planning, King Abdulaziz University (KAU), Saudi Arabia. He is an internationally recognized expert in urban and regional planning, quantitative methods, and the integration of digital technologies for sustainable development. He has demonstrated a proven ability to advance national and international priorities in urban sustainability, housing policy, and smart city development.

Manuscript received on 17 August 2025