




UNIVERSITÀ
DEGLI STUDI
DELL'AQUILA



Italiadomani
PIANO NAZIONALE
DI RIPRESA E RESILIENZA

**Dipartimento di Ingegneria e Scienze dell'Informazione
e Matematica**

PhD Thesis in Information and Communication Technologies (ICT)

Federated and Intelligent Digital Twins for Cultural Heritage: Frameworks, Applications, and Maturity Assessment

Supervisor

Prof. Henry Muccini

Co-Supervisor

Dr. Francesco Tarquini

Candidate

Gizealew Alazie Dagnaw

Student ID 289883

Anno Accademico 2025-2026

*From Digital Twins that mirrored
reality,
to Federated Twins that respected
autonomy,
to Cognitive Federated Digital Twins
that learned, reasoned, and adapted*

Declaration

I hereby declare that the work presented in this dissertation is my original research, carried out independently and in collaboration with my supervisors where explicitly acknowledged. This dissertation, entitled “*Federated and Intelligent Digital Twins for Cultural Heritage: Frameworks, Applications, and Maturity Assessment*” , has not been submitted, in whole or in part, for any other academic degree or qualification at this or any other institution.

All sources of information, data, figures, and ideas derived from the work of others have been properly cited and referenced in accordance with academic standards. Any collaborative contributions, published or under review manuscripts, and previously disseminated work forming part of this thesis are clearly indicated and appropriately acknowledged.

I affirm that this research was conducted ethically and responsibly, adhering to the institutional, legal, and professional guidelines governing research integrity, data protection, and scholarly conduct.

Gizealew Alazie Dagnaw

Place L'Aquila, Italy

Date:15/06/2026

Abstract

Context— Cultural heritage institutions face increasing challenges related to environmental degradation, structural vulnerabilities, visitor congestion, and resource limitations. At the same time, the digitalization of cultural assets has led to fragmented and siloed datasets distributed across heterogeneous systems. While Digital Twins (DTs) have the potential to support real-time monitoring, predictive simulation, and operational optimization, their application in cultural heritage remains constrained by issues of scalability, data sovereignty, and the absence of adaptive, multi-institution intelligence. Most existing approaches rely on static or centralized digital models that are insufficient for the dynamic, collaborative, and ethically sensitive nature of heritage management.

Objective— This dissertation aims to develop a comprehensive framework that advances DT technology from isolated digital replicas to adaptive, privacy-preserving, and cognitively enriched federated ecosystems. The research seeks to: (i) systematically evaluate the current state and maturity of DTs in cultural heritage; (ii) design multi-layered Federated and Cognitive Digital Twin architectures; and (iii) validate their ability to support decentralized intelligence, real-time emergency coordination, and knowledge-driven decision-making across multiple heritage institutions.

Method— The study employs a multi-stage methodology comprising: (1) a systematic review of 108 scholarly works to identify DT trends, enabling technologies, maturity levels, user-centered application and research gaps; (2) architectural design of federated frameworks, including an Intelligent FDT for distributed decision-making, and Cognitive Federated Digital Twin (CFDT) integrating semantic reasoning, reinforcement learning, and the MAPE-K loop; (3) simulation-based validation using agent-based modeling, federated learning (FedProx), and reinforcement learning with proximal policy optimization (PPO) in multi-museum evacuation scenarios; and (4) cross-framework synthesis and maturity assessment evaluating scalability, responsiveness, and cognitive adaptability.

Results— Findings from the systematic review highlight fragmented infrastructures, low DT maturity, and limited integration of AI, XR, or federated methods in heritage contexts. The proposed architectures demonstrate significant advancements: the Intelligent FDT framework improves emergency coordination with enhanced evacuation efficiency and reduced latency; and the CFDT provides stable learning, faster convergence, semantic explainability, and cross-institution knowledge transfer. Federated reinforcement learning yields higher cumulative rewards while reducing communication overhead and maintaining data sovereignty.

Conclusion— This dissertation establishes Intelligent, Federated, and Cognitive Digital Twin as a novel paradigm for cultural heritage management one that is adaptive, collaborative, ethically aligned, and scalable. By integrating DTs with AI, and federated intelligence, the work enables a transition from static digital archives to dynamic, learning ecosystems capable of supporting predictive conservation, real-time emergency response, and long-term sustainability. The resulting frameworks offer foundational theory, validated methods, and practical pathways for deploying intelligent federated digital infrastructures across the global cultural heritage sector.

Keywords: Digital Twin (DT), Federated Digital Twin (FDT), Cognitive Digital Twin (CDT), Federated Learning (FL), Cultural Heritage Management, Emergency Response Simulation, Agent-Based Modeling (ABM), Reinforcement Learning (RL), Privacy-Preserving Collaboration

Publication

This thesis is strongly based on the publications listed below:

1. Dagnaw, G., Capuano, R., Muccini, H. (2026). Digital twins for cultural heritage: A systematic analysis of the state of the art. *ACM Computing Surveys*, 58(9), Article 237.
2. Dagnaw, G. A., Muccini, H. (2026). AI-driven federated digital twins for scalable emergency response in cultural heritage sites. *ACM Journal on Computing and Cultural Heritage (JOCCH)*. (Accepted for publication, 2026)
3. Dagnaw, G. A., Muccini, H. (2026). Architecting Cognitive Federated Digital Twins with LLM Agents for Privacy-Preserving Adaptive Systems. *IEEE ICSA track 2026* (accepted)
4. Dagnaw, G. A., Martusciello, F. (2026). From Monitoring to Immersive Experiences: A Federated Digital Twin Framework for Intelligent Cultural Heritage Ecosystems, *Heritechne 2026 International Conference* (accepted)
5. Dagnaw, G. A., Muccini, H. (2025). Federated digital twins and XR for cultural heritage: A collaborative framework. In *Proceedings of the IMEKO MetroArchaeo 2025 Conference (IMEKO TC26 International Conference on Metrology for Archaeology and Cultural Heritage)*, Bergamo, Italy.
6. Dagnaw, G. A., Muccini, H. (2025). AI-powered federated digital twins for enhanced cultural heritage preservation and predictive analytics. In *ERCIM News*, Issue 141
7. Dagnaw, G. A., Muccini, H. (2025). Intelligent federated digital twin for cultural heritage: Collaborative preservation and visitor experience. *11th Italian Conference on ICT for Smart Cities and Communities*
8. Dagnaw, G. A., Muccini, H. (2026). Cognitive Federated Digital Twin Framework for Sustainable Cultural Heritage Preservation. *Information and Communication Technology for Sustainability (ICT4S 2026)* (submitted)
9. Dagnaw, G. A., Muccini, H. (2025). Decentralized 3D Heritage Digitization with AI-Enhanced Federated Digital Twins for 3D Research Challenges in Cultural Heritage – Vol. VI: AI-Driven 3D Digitisation Modelling in Cultural Heritage (submitted as book chapter).

Contents

List of Figures	viii
List of Tables	x
1 Introduction	1
1.1 Problem Statement	2
1.2 Research Motivation	4
1.3 Research Objectives	5
1.4 Research Questions	5
1.5 Methodological Overview	6
1.6 Significance of the Study	6
1.7 Contribution Summary	7
1.8 Thesis Structure	8
2 Background	11
2.1 Definition and Evolution of Digital Twins	11
2.1.1 Industrial Applications of Digital Twins	12
2.1.2 Stakeholder Requirements for Intelligent Digital Twin Adoption in Cultural Heritage	12
2.1.3 Application of Digital Twins in Cultural Heritage	13
2.1.4 Advantages and Challenges in Heritage Context	13
2.1.5 Motivation for Federated Learning in Heritage	14
2.2 Intelligent and Federated Digital Twins in Cultural Heritage	14
2.2.1 Applications of FDTs for Cultural Heritage	14
2.2.2 Cognitive Evolution of Digital Twins	15
2.3 Digital Twin Maturity and Adaptation in Heritage	15
2.4 Integration of DT, FL, and AI in Heritage Management	16
2.5 Summary of Limitations and Research Opportunities	16
3 Systematic Analysis of Digital Twin Applications in Cultural Heritage	17
3.1 Introduction	17
3.1.1 Motivation and Paper Contributions	18
3.1.2 Paper Structure	19
3.2 Background and Related Work	20
3.2.1 Digital Twin for Cultural Heritage	20
3.2.2 Related Work on Digital Twin for Cultural Heritage	21
3.3 Research Method	24
3.3.1 Planning the Review	24
3.3.2 Conducting and Reporting the Review	27
3.4 Digital Twin Technologies for Tangible Cultural Heritage Management (RQ1)	27
3.4.1 Adoption of Digital Twin Technologies (RQ1.1)	28
3.4.2 Immersive Technologies for Improving Stakeholder Experience (RQ1.2). . . .	30

3.4.3	Requirements of the Multidisciplinary User Community	33
3.5	Technology Usage of Digital Twin in Tangible Cultural Heritage (RQ2)	34
3.5.1	Tools and Platforms used for the Deployment of Digital Twin (RQ2.1)	34
3.5.2	Digital Twin Data Acquisition Techniques (RQ2.2)	36
3.5.3	Digital Twin Intelligence (RQ2.3)	38
3.6	The Maturity Level of Digital Twin in Tangible Cultural Heritage (RQ3)	40
3.6.1	Models to Assess the Maturity Level of Digital Twin (RQ3.1)	40
3.6.2	Implementation of Digital Twin (RQ3.2)	43
3.7	Discussion	46
3.7.1	Historical context and evolution DT	49
3.8	Open Challenges and Future Directions	50
3.8.1	Open Challenges from the Analysis	50
3.8.2	Future Directions	51
3.9	Conclusion	52
4	AI-Driven Federated Digital Twin for Cultural Heritage	53
4.1	Introduction	53
4.2	Background	55
4.2.1	Standalone Digital Twin vs Federated Digital Twin	55
4.2.2	Related Work	56
4.3	Proposed Federated Digital Twin Framework	57
4.3.1	Simulation-Based Validation	60
4.3.2	Empirical Evaluation of the Federated Digital Twin Architecture through Sim- ulation Results	62
4.4	Discussion	65
4.4.1	Case Study Scenario: Coordinated Emergency Response Across Two Museums	66
4.4.2	Limitations and Modeling Realism	67
4.4.3	Toward Real-World Deployment	67
4.5	Conclusions and Future Work	67
5	Real-Time Coordination in Federated Digital Twins	70
5.1	Introduction	70
5.2	Background	72
5.2.1	Federated Digital Twins (FDTs)	73
5.2.2	Real-Time Systems and WCET Modeling	73
5.3	Related Work	74
5.4	Federated Real-Time Embedded Digital Twin Integration Framework	74
5.4.1	Federated Learning Integration and Coordination	76
5.4.2	Mathematical Model and Real-Time Guarantees	77
5.5	Simulation Setup and Evaluation	79
5.5.1	Rationale for Simulation-Based Evaluation	79
5.5.2	Real-Time Evaluation and WCET Profiling	80
5.5.3	Task-Level Execution and Cycle Profiling	81
5.6	Simulation Outcomes and Performance Analysis	82
5.6.1	Extended Results and Comparative Evaluation	83
5.6.2	Comparative Performance Analysis	84
5.6.3	Scalability Evaluation under Increasing Federation Size	84
5.6.4	Ablation Study on System Components	85
5.6.5	Interpretation and Industrial Relevance	85
5.7	Discussion and Implications	85
5.7.1	Theoretical and Architectural Implications	86
5.7.2	Practical and Industrial Implications	86

5.7.3	Integration with Industrial Information Systems	86
5.7.4	Societal and Sustainability Implications	87
5.7.5	Limitations and Future Directions	87
5.7.6	Broader Research Implications	87
5.8	Conclusions and Future Work	87
6	Cognitive Federated Digital Twins for Sustainable Cultural Heritage	90
6.1	Introduction	91
6.1.1	Research Questions	92
6.1.2	Contributions	92
6.2	Background	93
6.2.1	Digital Twin Technology in Cultural Heritage	93
6.2.2	Federated Learning for Distributed Knowledge Sharing	94
6.2.3	Cognitive and Self-Learning Digital Twins	94
6.2.4	The MAPE-K Loop for Self-Adaptation	94
6.2.5	Research Gap and Contribution	95
6.3	Methodology	95
6.3.1	Overview of the Cognitive Federated Digital Twin (CFDT) Framework	95
6.3.2	System Architecture and Components	95
6.3.3	MAPE-K Driven Adaptive Reinforcement Learning Loop	97
6.3.4	Federated Reinforcement Learning with FedProx	97
6.3.5	Reward Function	98
6.3.6	Reinforcement Learning and PPO Policy Optimization	98
6.3.7	Digital Twin Simulation and Adaptation Logic	98
6.3.8	Baseline Methods	99
6.3.9	Simulation Environment and Execution Protocol	99
6.3.10	Algorithm CFDT Training Process	100
6.4	Findings	101
6.4.1	Quantitative findings	101
6.4.2	Qualitative Findings	101
6.5	Results and Discussion	102
6.5.1	Overview of Simulation Outcomes	102
6.5.2	Convergence and Learning Behavior Analysis	102
6.5.3	Role of the MAPE-K Loop in System Adaptation	103
6.5.4	Inter-Museum Collaborative Learning Dynamics	104
6.5.5	Analysis of Communication Efficiency and Latency	105
6.5.6	Behavior of Multi-Objective Optimization	106
6.5.7	Comparative Discussion with Prior Research	106
6.5.8	Interpretation of MAPE-K Evolution and Cognitive Maturity	107
6.5.9	Scalability and Generalization Implications	107
6.5.10	Resilience and Stability of the Federated Digital Twin System	107
6.5.11	Interpretation of Final Knowledge State	108
6.5.12	Summary of Insights and Implications	108
6.6	Conclusion and Future Work	109
7	Discussion and Synthesis	111
7.1	Synthesis Across Research Stages	111
7.1.1	From Systematic Review to Conceptualization	111
7.1.2	Intelligent Federated Digital Twin Framework	112
7.1.3	Real-Time Federated Digital Twin Framework	113
7.1.4	Cognitive Federated Digital Twin Framework	113
7.2	Theoretical Implications	114

7.3	Ethical and Cultural Dimensions	114
7.3.1	Data Quality, Provenance, and FAIR/CARE Compliance	115
7.4	Interoperability and Policy Implications	115
8	Conclusion and Future Work	117
8.1	Scientific and Practical Implications	118
8.2	Limitations	118
8.3	Future Work	118
8.3.1	Technical Enhancements and AI Integration	118
8.3.2	Immersive Interfaces and Spatial Realism	119
8.3.3	Real-World Deployment and Validation	119
8.3.4	Ethical Governance and Open Science	119
8.3.5	Scalability and Sustainability	119
	bibliography	121

List of Figures

1.1	Thesis roadmap and research flow.	10
3.1	Screening Phase Description and Numbers	26
3.2	Taxonomy-informed conceptual framework for <i>DT</i> applications in tangible cultural heritage, structured along three complementary dimensions: user-centric applications, enabling technologies, and maturity levels, along with their interdependencies. Arrow thickness indicates relationship strength: thick lines = strong, medium lines = moderate, thin lines = weak.	48
3.3	Timeline of <i>DT</i> evolution in cultural heritage (2002–2025), highlighting milestones from conceptual origin to intelligent and federated <i>DT</i> frameworks. Adapted from [7, 10, 51, 87, 97, 101, 122, 123, 169, 239].	50
4.1	Comparison of Digital Twin and Federated Digital Twin across key technical and operational aspects.	56
4.2	Multi-layered framework for federated digital twin management in museums.	59
4.3	Initial agent distribution and hazard zones.	65
4.4	Mid-simulation showing evacuation flow.	65
4.5	Aggregated performance metrics across runs.	66
5.1	Evolution of Digital Twin to Federated Digital Twins, showing major technological and conceptual milestones from NASA’s early developments in the 1970s to the expansion of Federated Digital Twin research in 2024.	71
5.2	Federated Real-Time Embedded Digital Twin Integration Framework (FRT-DTIF) for Emergency Coordination.	76
5.3	Real-time Federated Digital Twin Control Loop and WCET Profiling.	82
5.4	Simulation dashboard showing WCET cycle time, synchronization latency, and evacuation performance.	83
6.1	Proposed Cognitive Federated Digital Twin (CFDT) architecture for interconnected museums. Local DTs interact with physical layers and reinforcement learning agents, while the cognitive MAPE-K loop enables adaptation. A secure federated layer aggregates knowledge without exposing private data.	96
6.2	Performance comparison between Local-Only RL (Museums A–D) and the proposed CFDT framework. The CFDT system achieves higher cumulative rewards and smoother convergence, demonstrating improved learning stability and cross-museum knowledge transfer.	103
6.3	Federated reward convergence across 10 rounds for the CFDT system.	104
6.4	Coordination efficiency trends across federated rounds. The CFDT framework maintains high coordination efficiency ($\approx 3550\% - 3850\%$) demonstrating stable collaborative learning across distributed museum environments.	105
6.5	Communication volume vs. system efficiency across federated rounds.	106
6.6	Federated learning reward progression across Museums A–D using PPO + FedProx integrated with NetLogo simulation.	109

6.7	Floor plans for the four simulated museums used in the NetLogo environment. Each museum contains three rooms and a central exhibition space: Museum A (red), Museum B (blue), Museum C (green), Museum D (yellow).	110
-----	--	-----

List of Tables

3.1	Comparative analysis of prior surveys with respect to the three dimensions considered in this study	23
3.2	Research Questions	24
3.3	Search Keywords: Terms and Synonyms	25
3.4	Inclusion and exclusion criteria used	26
3.5	Selected Parameters for the Study	27
3.6	Co-occurrence of studies on types of artworks (<i>RQ1.1.P1</i>) and Purpose of the CH digital twin (<i>RQ1.1.P2</i>).	29
3.7	Co-occurrence of Purpose of the TCH digital twin (<i>RQ1.1.P2</i>) and digital twin techniques (<i>RQ1.1.P3</i>).	31
3.8	Co-occurrence of studies on access devices (<i>RQ1.2.P2</i>) and end users of digital heritage (<i>RQ1.2.P3</i>).	33
3.9	Co-occurrence of studies on Access devices (<i>RQ1.2.P2</i>) and immersive technologies (<i>RQ1.2.P1</i>).	34
3.10	Co-occurrence of studies on data-acquisition technology (<i>RQ2.2.P1</i>) and Type of data (<i>RQ2.2.P2</i>).	37
3.11	Co-occurrence of studies on activity level (<i>RQ2.3.P1</i>) and functional layer of digital twin (<i>RQ2.3.P2</i>).	39
3.12	Co-occurrence of studies on model update frequency (<i>RQ3.1.P1</i>) and twin models (<i>RQ3.1.P2</i>).	43
3.13	Co-occurrence of studies on model update frequency (<i>RQ3.1.P1</i>) and Model security (<i>RQ3.1.P3</i>).	44
3.14	Co-occurrence of studies on model update frequency (<i>RQ3.1.P1</i>) and Verification and validation of DT models (<i>RQ3.1.P4</i>).	45
3.15	Co-occurrence of studies on stage of digital twin implementation (<i>RQ3.2.P1</i>) and connectivity mode (<i>RQ3.2.P2</i>).	46
3.16	Co-occurrence of studies on stage of digital twin implementation (<i>RQ3.2.P1</i>) and Computational intelligence (<i>RQ3.2.P3</i>).	47
4.1	Summary of Key Simulation Metrics, Definitions, Benchmarks, and Results for Federated Digital Twin Evaluation	64
4.2	Federated Learning Hyperparameters and Model Configuration	64
5.1	Comparative Analysis of Related Work	75
5.2	Cycle Performance Summary (1120-Tick Simulation)	81
5.3	Comparative performance of FRT-DTIF versus baseline systems (mean $\pm$ 95% CI)	84
5.4	Ablation study of system components	85
6.1	Hyperparameters for CFDT Simulation	100
6.2	Metric definitions and computation method	100
6.3	Summary of Federated Learning Results across Ten Rounds	103

Chapter 1

Introduction

The preservation and management of cultural heritage have long been recognized as essential to the continuity of human civilization [36]. Cultural heritage encompasses both tangible and intangible expressions of humanity’s collective memory, including monuments, museums, archaeological sites, artifacts, oral traditions, and cultural landscapes [131]. As cultural resources, these assets embody the social, historical, and artistic achievements of past generations and provide the foundation for identity formation, education, and intercultural dialogue [243]. However, in the twenty-first century, cultural heritage faces an expanding array of threats ranging from environmental degradation and natural disasters to armed conflict, uncontrolled urbanization, and insufficient conservation funding. According to UNESCO’s 2022 Global Heritage Outlook, more than one-third of the world’s recognized heritage sites are exposed to climate-related hazards such as flooding, heat stress, and sea-level rise [251]. These circumstances demand a paradigm shift in how heritage is documented, protected, and transmitted to future generations.

In parallel, the advent of the digital transformation has introduced new opportunities for heritage conservation, accessibility, and public engagement. Over the last two decades, digital technologies such as 3D laser scanning, photogrammetry, Building Information Modelling (BIM), the Internet of Things (IoT), and artificial intelligence (AI) have transformed the methods by which cultural data are captured, stored, and analyzed [228]. These tools enable the creation of high-fidelity virtual representations of artifacts and heritage environments, supporting research, restoration, and dissemination. Moreover, immersive media such as Virtual Reality (VR), Augmented Reality (AR), and Mixed Reality (MR), collectively referred to as Extended Reality (XR) are redefining visitor experiences by enabling interactive storytelling and multisensory engagement [18, 42]. The European Commission’s *Digital Heritage Strategy 2030* underscores that digitalization is not merely an auxiliary tool but a transformative domain that must be embedded into the entire heritage lifecycle from documentation to interpretation and access [29, 36, 64, 168].

Nevertheless, the growing complexity of digital heritage ecosystems introduces new challenges of scalability, data governance, and interoperability [149]. Large cultural institutions generate vast and heterogeneous datasets from high-resolution scans to visitor interaction logs often stored in isolated silos, governed by different technical standards, and managed by distinct authorities. This fragmentation impedes data sharing and cross-institutional collaboration [121]. Furthermore, most current digital heritage systems are static in nature, they produce digital replicas of artifacts or buildings that lack real-time synchronization with the physical counterparts. This static representation limits

predictive analysis and adaptive management, especially in dynamic operational contexts such as museums or heritage cities [18].

The *Digital Twin* (DT) concept offers a compelling pathway to overcome the above limitations. Originating in industrial engineering and aerospace domains, the digital twin paradigm describes a real-time, data-driven digital replica of a physical asset, system, or process, capable of bidirectional communication and continuous learning [241]. Through the integration of IoT sensors, machine learning algorithms, DTs enable continuous monitoring, predictive maintenance, and performance optimization in complex systems [86]. Over time, the DT concept has evolved beyond engineering to domains such as healthcare, urban management, and environmental monitoring [117, 211]. In cultural heritage, DTs can serve as living models of museums, monuments, or archaeological sites that integrate multimodal data sources structural health sensors, environmental monitors, visitor tracking systems and feed them into intelligent decision support systems [30]. Such an approach provides both a preservation mechanism and a platform for enhanced engagement through interactive visualization.

However, as cultural heritage environments typically involve multiple stakeholders (museums, archives, local communities, governmental agencies, and international organizations), the deployment of a centralized digital twin architecture raises concerns about scalability, privacy, and institutional autonomy [151, 277]. Many heritage institutions operate under strict policies that prevent the external sharing of raw cultural data due to issues of ownership, authenticity, and security [69, 214]. To address these issues, researchers have begun exploring *Federated Digital Twin* (FDT) framework, where data and computational processes are distributed across multiple nodes that co-operate through federated learning (FL) mechanisms [139, 280]. In this paradigm, each institution maintains its own local digital twin, which contributes to a shared global intelligence model without transferring sensitive raw data.

Federated learning, initially proposed in the context of distributed artificial intelligence, allows multiple clients to collaboratively train a global model while preserving data privacy and locality [172]. When applied to digital twin systems, this approach fosters key dimensions of interoperability, scalability, and ethical compliance for multi-institutional cultural heritage management [290, 294]. It also supports adaptive behavior in real-time, local twins can autonomously react to environmental changes while benefiting from collective insights generated by the federation. For example, in an emergency such as a fire or structural collapse, federated digital twins can coordinate evacuation strategies across multiple museums or heritage buildings while respecting each institution's operational sovereignty [207].

1.1 Problem Statement

Despite the rapid digitization of cultural heritage assets, the transition from static digital documentation toward intelligent, adaptive, and collaborative ecosystems remains unexplored. Current digital heritage initiatives predominantly rely on fragmented infrastructures, isolated databases, and visualization tools that lack interoperability across institutions or geographic regions [141, 228]. These systems, though valuable for archival purposes, fail to provide real-time monitoring, predictive analysis, or self-adaptive functionality. Consequently, heritage institutions continue to operate reactively rather than proactively in addressing environmental degradation, structural stress, or

visitor-related risks [39].

The adoption of Digital Twin (DT) technology in cultural heritage remains constrained by a combination of technical, organizational, and ethical barriers. First, the absence of standardized data models and semantic interoperability frameworks hinders the integration of heterogeneous datasets derived from sensors, 3D scans, and historical archives. Unifying such diverse data ranging from geometric models to textual metadata governed by ontologies such as CIDOC CRM into a continuously updating digital twin remains a significant challenge [50, 121]. Second, limited real-time sensing and data fusion capabilities impede the creation of dynamic, responsive digital replicas. Most heritage DTs remain static, serving as digital surrogates for visualization or education rather than as living systems that support predictive or prescriptive decision-making [30]. Third, persistent concerns over data privacy, intellectual property, and ethical stewardship continue to discourage inter-institutional data sharing. Many organizations, particularly in the Global South, fear cultural misappropriation or misuse of sensitive information [64].

A further gap lies in the insufficient integration of Artificial Intelligence (AI) and autonomous learning mechanisms within heritage-oriented DT frameworks. While AI has revolutionized industries such as manufacturing, healthcare, and logistics, its role in cultural heritage remains exploratory and fragmented [293]. The lack of AI-driven intelligence inhibits real-time anomaly detection, predictive maintenance, and adaptive environmental control. Likewise, decision-making processes during emergencies remain largely manual, reactive, and centralized, limiting situational awareness and response efficiency.

Emergency management represents one of the most critical and underdeveloped areas in digital heritage. High-profile disasters such as the 2019 Notre-Dame Cathedral fire in Paris and the 2020 National Museum of Brazil blaze revealed the vulnerability of heritage sites to natural and human-induced hazards [85, 179]. Traditional emergency protocols depend on static evacuation maps, limited simulations, and delayed coordination, which are insufficient in rapidly evolving crisis contexts. Without adaptive intelligence and multi-site collaboration, institutions struggle to safeguard both human life and cultural assets.

The emerging paradigm of *Federated Digital Twins (FDTs)* offers a promising foundation for addressing these limitations. Through Federated Learning (FL), multiple heritage institutions can collaboratively train AI models without sharing sensitive data, thereby preserving cultural sovereignty and privacy [139, 280]. Each local twin maintains autonomy while contributing to a global intelligence layer that enhances system-wide situational awareness. In emergency scenarios, this approach enables decentralized yet coordinated decision-making, where each site can adapt its strategies based on local conditions while benefiting from collective insights across the federation [207].

However, even federated architectures remain inherently limited in their adaptability and reasoning capacity. They depend on predefined learning objectives and lack the ability to evolve autonomously in complex, uncertain environments. This dissertation therefore advances the concept of the *Cognitive Federated Digital Twin (CFDT)*, which integrates cognitive computing, semantic reasoning, and reinforcement learning into the federated structure. The CFDT model enables self-learning, contextual understanding, and ethically guided decision-making across distributed digital twins. By combining the privacy-preserving collaboration of FDTs with the self-awareness and reasoning capabilities of cognitive systems, CFDTs establish a new paradigm for intelligent, resilient, and ethically aligned cultural heritage management.

In this vision, each cognitive twin operates as a semi-autonomous agent within a federated ecosystem, capable of perceiving environmental changes, analyzing contextual data, and making adaptive decisions in real time. The cognitive layer, supported by the MAPE-K (Monitor–Analyze–Plan–Execute–Knowledge) feedback loop, continuously refines decision strategies based on feedback and learned experiences. This framework not only improves preservation and operational resilience but also embeds ethical reflection within autonomous processes, aligning technological interventions with cultural and societal values [69, 251].

Nonetheless, realizing such a Cognitive Federated Digital Twin ecosystem entails multiple unresolved challenges. The computational overhead of federated and cognitive learning may strain the limited digital infrastructure of smaller museums or community heritage centers [280]. Ensuring synchronization, model convergence, and semantic interoperability across distributed systems remains a complex research problem. Furthermore, current maturity assessment models are not equipped to evaluate the readiness of heritage institutions to adopt such advanced, multi-layered technologies [117]. These challenges underscore the need for a holistic framework that balances technical feasibility, ethical responsibility, and institutional inclusivity.

In this context, the central research problem addressed by this dissertation can be articulated as follows:

Our systematic literature review reveals a lack of intelligent, scalable, interoperable, and privacy-preserving frameworks that integrate cognitive intelligence and federated collaboration into Digital Twin systems for cultural heritage preservation, adaptive emergency response, and ethical decision-making.

This dissertation addresses this problem by proposing a comprehensive *multi-layered Federated and Cognitive Digital Twin (CFDT)* framework tailored to cultural heritage management. The framework emphasizes semantic interoperability, distributed intelligence, and ethical governance. It seeks to demonstrate how CFDTs can transform heritage institutions from static repositories of data into dynamic, learning ecosystems capable of collective intelligence, adaptive resilience, and culturally aligned decision-making.

1.2 Research Motivation

The motivation for this research arises from the urgent need to develop resilient, adaptive, and ethically governed digital infrastructures for the management and preservation of cultural heritage. Despite significant advances in digitization, most heritage systems remain static, reactive, and fragmented, lacking the ability to learn, adapt, or collaborate effectively across institutional and national boundaries. As the digital transformation of heritage accelerates, there is a growing demand for frameworks that can integrate distributed intelligence, ensure cultural data sovereignty, and support proactive, self-adaptive decision-making.

This dissertation is driven by the recognition that the convergence of *Artificial Intelligence (AI)*, *Federated Learning (FL)*, and cognitive computing within the *Digital Twin (DT)* ecosystem can revolutionize how cultural assets are preserved, monitored, and managed. Federated Learning enables distributed AI training across multiple institutions without exchanging sensitive cultural data, thereby ensuring privacy and respecting ethical and geopolitical boundaries [139, 172]. Cog-

nitive computing, when embedded within Digital Twins, enhances system intelligence by enabling reasoning, learning, and contextual understanding [154, 293].

This evolution gives rise to the *Cognitive Federated Digital Twin (CFDT)* paradigm a new generation of digital systems capable of collective intelligence and self-optimization across distributed heritage environments. CFDTs extend conventional FDTs by incorporating semantic reasoning, reinforcement learning, and feedback-driven adaptation, allowing heritage institutions to collaborate securely while autonomously improving their operations. In this way, CFDTs represent a shift from descriptive and predictive digital systems toward prescriptive and cognitive digital ecosystems that are not only technologically advanced but also ethically and culturally aligned.

By pursuing this line of inquiry, the research contributes to a future where cultural heritage management evolves from isolated digital replicas to interconnected cognitive ecosystems capable of safeguarding artifacts, predicting environmental risks, and supporting real-time, ethically grounded decision-making.

1.3 Research Objectives

The overarching goal of this dissertation is to design, implement, and validate a scalable, privacy-preserving, and cognitively enhanced *Federated Digital Twin Framework* for cultural heritage preservation, emergency response, and intelligent decision-making. Specifically, the research pursues the following objectives:

1. To systematically analyze the current state of Digital Twin applications in cultural heritage, identifying their maturity levels, enabling technologies, and user-centric applications.
2. To design a multi-layered Federated and Cognitive Digital Twin (CFDT) architecture that integrates AI, Federated Learning, and cognitive reasoning mechanisms for adaptive collaboration among heritage institutions.
3. To develop intelligent learning components that enable decentralized, self-adaptive decision-making in scenarios such as emergency management, structural monitoring, and visitor flow optimization.
4. To validate the proposed framework through agent-based simulation and federated learning experiments, assessing its scalability, performance, and ethical robustness.

1.4 Research Questions

To achieve the above objectives, this study addresses the following research questions:

- How can Digital Twin technologies evolve from static representations to cognitive, federated ecosystems that support adaptive cultural heritage management?
- What architectural and algorithmic mechanisms enable the integration of Federated Learning and cognitive reasoning into Digital Twin frameworks while ensuring privacy, scalability, and semantic interoperability?

- How can AI-driven learning models and feedback loops (e.g., MAPE-K) enhance the decision-making capacity of heritage institutions under uncertain and dynamic conditions such as emergencies?
- What simulation-based and federated learning methods can effectively validate the performance, adaptability, and ethical behavior of Cognitive Federated Digital Twin systems?

1.5 Methodological Overview

This research employs a mixed-methods approach combining systematic literature analysis, architectural design, simulation-based experimentation, and maturity assessment. The methodology is structured around four interrelated phases aligned with the dissertation’s major research outputs:

1. **Systematic Review:** A comprehensive review of Digital Twin and Federated Learning applications in cultural heritage, following Kitchenham and PRISMA guidelines to identify trends, challenges, and gaps.
2. **Framework Design:** Development of multi-layered Federated and Cognitive Digital Twin (CFDT) architecture that integrates perception, local twin, federated coordination, and cognitive intelligence layers.
3. **Simulation and Validation:** Implementation of agent-based and federated learning models to simulate dynamic scenarios (e.g., multi-museum emergency coordination) and evaluate performance metrics such as scalability, latency, and decision accuracy.
4. **Maturity and Ethical Assessment:** Formulation of a cognitive maturity model that captures the technical, operational, and ethical dimensions of CFDT deployment across heterogeneous heritage ecosystems.

1.6 Significance of the Study

The significance of this research lies in its contribution to advancing both the theoretical foundations and the practical realization of intelligent, privacy-preserving digital ecosystems for cultural heritage. Theoretically, the study extends the concept of the *Digital Twin (DT)* by introducing Federated and Cognitive intelligence, enabling autonomous learning, reasoning, and adaptive control within distributed heritage environments. It redefines the traditional DT model by embedding self-awareness, contextual reasoning, and ethical decision-making core attributes that elevate digital systems from reactive simulations to proactive, cognitive entities [154, 293].

Technically, the research introduces a scalable *a multi-layered Federated and Cognitive Digital Twin (CFDT)* architecture that integrates Artificial Intelligence (AI), Federated Learning (FL), and semantic reasoning to enable secure, multi-institutional collaboration. This architecture allows heritage institutions to jointly train intelligent models without sharing sensitive cultural data, ensuring data sovereignty and compliance with ethical governance principles. The incorporation of cognitive mechanisms such as the MAPE-K feedback loop and knowledge-based reasoning empowers these systems to adapt autonomously to environmental changes, visitor dynamics, and operational challenges.

From a societal and policy perspective, this research contributes to the sustainable and inclusive management of cultural heritage. By enabling decentralized intelligence and participatory collaboration, the CFDT paradigm supports institutions to improve preservation, operational efficiency, and public participation while maintaining cultural integrity and ethical accountability. These contributions align with the United Nations Sustainable Development Goals (SDGs), particularly: SDG 11: Sustainable Cities and Communities through the preservation and digital stewardship of cultural heritage assets using intelligent, data-driven systems. In general, this dissertation offers a new theoretical and practical framework for transforming heritage management from static documentation into a dynamic, cognitive, and federated ecosystem that balances innovation, ethics, and sustainability.

1.7 Contribution Summary

This dissertation makes several original and substantive research contributions that advance the state of the art in Digital Twins (DT), Federated Digital Twins (FDT), and Cognitive Federated Digital Twins (CFDT) for cultural heritage. The main contributions are summarized as follows:

- **Comprehensive Systematic Literature Review:** Conducted an extensive systematic analysis of 108 studies on Digital Twin applications in cultural heritage, developing a tri-dimensional framework covering user-centric applications, enabling technologies, and maturity levels. This review identified critical gaps including low DT maturity, fragmented infrastructures, limited AI/FL/XR adoption, and a lack of interoperability.
- **Intelligent Federated Digital Twin (FDT) Framework:** Designed a multi-layered FDT architecture enabling privacy-preserving, decentralized collaboration across cultural heritage institutions. The architecture incorporates semantic interoperability, agent-based modeling, federated learning, and adaptive decision-making.
- **Real-Time Federated Digital Twin Framework:** Developed a Real-Time Federated Digital Twin Framework (RT-FDTF) featuring deterministic scheduling, Worst-Case Execution Time (WCET) profiling, and bounded synchronization latency. Provided mathematical modeling ensuring that all decision-making cycles execute within a 500 ms deadline.
- **Cognitive Federated Digital Twin (CFDT):** Introduced the Cognitive Federated Digital Twin, integrating reinforcement learning (PPO), federated reinforcement learning (FedProx), semantic reasoning, and the MAPE-K loop to enable self-adaptation, cross-museum knowledge sharing, and ethical decision-making.
- **Simulation-Based Validation:** Developed a hybrid simulation framework combining NetLogo agent-based modeling with federated learning to validate large-scale emergency scenarios across multiple museums. Experiments demonstrated high performance, including a 93.8% evacuation success rate across simulation runs, 88.8% routing efficiency, and strong federated convergence.
- **Cultural Heritage-Specific Maturity Model:** Proposed a four-level maturity model (Pre-DT, DT, Federated DT, Cognitive Federated DT) tailored to cultural heritage, incorporating ethical, societal, and operational dimensions not addressed by industrial models.

- **Ethical, Sustainable, and Societal Impact:** Embedded principles of ethical AI, cultural data sovereignty, GDPR-compliant practices, and UNESCO AI ethics. Contributions directly support sustainable and resilient cultural heritage management aligned with SDG 11.
- **Scientific Publications:** Supported by peer-reviewed publications, including ACM Computing Surveys, ACM Journal on Computing and Cultural Heritage, IEEE ICSCA , and other respected venues, demonstrating strong scholarly impact.

These contributions collectively advance digital twin research toward collaborative and intelligent cultural heritage ecosystems.

1.8 Thesis Structure

This dissertation is organized into eight interrelated chapters, each contributing to the overarching goal of developing and validating a comprehensive framework for Digital Twin evolution, progressing from conventional Digital Twins to Federated Digital Twins, and ultimately to Cognitive Federated Digital Twins (CFDTs), with a focus on cultural heritage preservation and emergency management. To provide a clearer conceptual understanding of the dissertation flow and the progressive evolution of the proposed frameworks, Figure 1.1 presents a graphical overview of the thesis structure. The figure illustrates the logical progression of the research from foundational concepts and theoretical analysis toward the development, validation, and cognitive enhancement of Federated Digital Twin frameworks for cultural heritage applications. Furthermore, it highlights the relationship between individual chapters, the associated research phases, and the principal scientific contributions achieved throughout the dissertation.

Chapter 1 introduces the overall research context, articulating the background, problem statement, motivation, research objectives, guiding questions, and methodological overview, as well as emphasizing the theoretical and societal significance of the work. Following this, **Chapter 2** provides a comprehensive review of the theoretical and technological foundations that underpin the study. It explores the evolution of Digital Twin (DT) technologies, Federated Learning (FL), and Cognitive Digital Twins, analyzing their potential applications and ethical implications within cultural heritage environments.

Chapter 3 presents a systematic analysis of existing Digital Twin applications in cultural heritage. This chapter synthesizes the results of a structured literature review to identify current trends, user-centered applications, enabling technologies, and maturity levels of DT adoption, as well as the critical gaps that motivate the proposed research framework.

Building on these insights, **Chapter 4** proposes an intelligent, AI-driven Federated Digital Twin framework for cultural heritage management. It incorporates adaptive decision-making mechanisms and validates the proposed architecture through agent-based simulations, demonstrating its capability for efficient, privacy-preserving collaboration.

Following this, **Chapter 5** presents a real-time, deterministic Federated Digital Twin framework designed for emergency response coordination. This chapter integrates Worst-Case Execution Time (WCET)-aware scheduling and synchronization mechanisms to enable time-critical collaboration between heritage sites and validates these capabilities through extensive simulation-based evaluation.

Chapter 6 introduces the Cognitive Federated Digital Twin (CFDT) model the advanced contribution of this dissertation. This framework unifies federated reinforcement learning, semantic reasoning, and the MAPE-K (Monitor–Analyze–Plan–Execute over Knowledge) control loop to enable adaptive, knowledge-driven decision-making across distributed heritage ecosystems. The CFDT architecture represents a paradigm shift toward intelligent, self-learning heritage management systems capable of dynamic adaptation to environmental, structural, and human factors.

Chapter 7 synthesizes the findings derived from all preceding research stages, critically examining their theoretical, technical, and practical implications. It also discusses cross-chapter interconnections and positions the CFDT framework within broader discourses of digital heritage, artificial intelligence, and sustainable technological transformation. Finally, **Chapter 8** concludes the dissertation by summarizing its major contributions, outlining theoretical advancements, identifying limitations, and proposing potential directions for future work in the development of intelligent and federated digital ecosystems for cultural heritage.



Figure 1.1: Thesis roadmap and research flow.

Chapter 2

Background

Digital Twins, originally developed for industrial applications, provide a real-time, bidirectional connection between physical assets and their virtual counterparts, enabling continuous monitoring, predictive analysis, and decision-making [154, 241]. When applied to cultural heritage, DTs offer the potential to dynamically simulate environmental conditions, visitor behavior, and conservation processes, thereby facilitating more effective preservation strategies and improved visitor experiences [18, 30]. Furthermore, Federated Learning offers a framework for collaborative intelligence across multiple heritage institutions while preserving data privacy, a critical requirement due to the sensitive nature of cultural data [139, 172, 280].

This chapter provides a comprehensive review of the theoretical foundations and related work underpinning this research. It begins by tracing the evolution of Digital Twin technology and its applications in both industrial and cultural heritage contexts. It then explores Federated Learning and privacy-preserving AI techniques, emphasizing their relevance to multi-institutional heritage management. The concept of Cognitive Federated Digital Twins (CFDTs) is subsequently introduced and discussed. Furthermore, existing Digital Twin maturity models are reviewed, with a focus on their limitations in cultural heritage scenarios. Finally, the chapter identifies gaps in the current literature and outlines the research opportunities that motivate this dissertation.

2.1 Definition and Evolution of Digital Twins

Digital Twins are virtual representations of physical systems that are continuously updated with real-time data collected from sensors, instruments, and historical records [241]. The concept was first introduced within industrial settings, where DTs were used to monitor complex machinery, optimize operations, and predict failures [241]. Over time, the technology has evolved through several stages. Initially, descriptive twins provided static digital replicas of physical assets. Diagnostic twins incorporated sensor data to monitor real-time system behavior and detect anomalies. Predictive twins further employed statistical and machine learning algorithms to forecast future states. The most advanced form, intelligent twins, combines predictive capabilities with prescriptive and autonomous decision-making, enabling systems to self-optimize and adapt to changing conditions [117, 293].

The evolution of DTs has been facilitated by advances in sensor technology, IoT networks, cloud computing, and AI-driven analytics [154]. In particular, AI allows for the extraction of actionable

insights from large volumes of heterogeneous data, supporting predictive maintenance, resource optimization, and decision-making processes. The continuous interaction between the physical asset and its digital counterpart forms the foundation of the DT paradigm, offering unprecedented opportunities for real-time monitoring, simulation, and management [117, 238].

2.1.1 Industrial Applications of Digital Twins

DTs have been widely adopted in industries such as aerospace, automotive manufacturing, and smart factories, where they support predictive maintenance, operational optimization, and system design improvements [117]. In aerospace, for instance, DTs monitor aircraft components during flight and predict potential failures, enabling preemptive maintenance that minimizes downtime and enhances safety. In manufacturing, DTs facilitate real-time monitoring of production lines, simulation of operational scenarios, and optimization of resource allocation. These industrial applications highlight the potential of DTs to improve efficiency, reduce costs, and enhance system reliability [154, 241].

2.1.2 Stakeholder Requirements for Intelligent Digital Twin Adoption in Cultural Heritage

While Digital Twins (DTs) have demonstrated significant value in industrial domains through predictive maintenance, operational optimization, and intelligent decision support, their adoption within Cultural Heritage (CH) environments is driven by a broader set of operational, cultural, and societal requirements [79, 113]. Unlike industrial systems, where efficiency and productivity are often the primary objectives, cultural heritage ecosystems involve diverse stakeholders whose priorities extend beyond technical performance to include preservation, accessibility, education, sustainability, and cultural continuity.

The development of Digital Twins for Cultural Heritage is therefore not solely motivated by technological innovation but also by the strategic and operational needs of multiple stakeholders involved throughout the lifecycle of heritage assets [79, 113]. While existing DT implementations and Historic Building Information Modeling (HBIM) workflows primarily focus on geometric representation, continuous monitoring, and localized simulation functionalities [152], the long-term sustainability and usefulness of DTs strongly depend on their ability to integrate heterogeneous knowledge resources within a unified and semantically coherent framework [187].

From the perspective of heritage owners and custodians, Digital Twins provide several strategic advantages, including lifecycle documentation, intervention planning, preventive conservation, risk monitoring, sustainability assessment, and decision support for heritage management activities [71]. Furthermore, DTs facilitate evidence-based preservation policies by enabling the continuous integration of monitoring data, historical records, and contextual heritage information into a comprehensive digital environment [79].

In addition, stakeholders such as conservators, archaeologists, curators, educators, policymakers, tourists, and local communities require distinct forms of interaction with heritage data and digital services [107]. Consequently, Digital Twins must evolve beyond purely technical replicas toward intelligent, knowledge-centric ecosystems capable of supporting interdisciplinary collaboration, knowledge sharing, and long-term cultural memory preservation [71].

To satisfy these diverse stakeholder requirements, DT ecosystems must incorporate not only

sensor-generated data but also contextual narratives, semantic relationships, provenance metadata, and intangible cultural knowledge. The integration of these heterogeneous information sources requires semantic interoperability mechanisms based on standardized cultural heritage ontologies, such as the CIDOC Conceptual Reference Model (CIDOC CRM) [68]. Such semantic foundations enable the representation, exchange, and interpretation of knowledge across institutions, disciplines, and technological platforms.

Furthermore, the justification for developing intelligent and federated Digital Twin infrastructures extends beyond computational capabilities alone [207]. These systems must also align with internationally recognized principles for data governance, stewardship, and ethical management. In particular, cultural heritage data should be maintained, shared, and evaluated according to the FAIR (Findable, Accessible, Interoperable, and Reusable) and CARE (Collective Benefit, Authority to Control, Responsibility, and Ethics) principles [37, 242]. Adherence to these frameworks ensures that Digital Twins evolve from isolated technological artifacts into trustworthy, sustainable, and ethically governed knowledge infrastructures capable of supporting collaborative cultural heritage ecosystems.

The stakeholder-driven requirements outlined above reveal that conventional Digital Twins are insufficient to address the complexity of contemporary cultural heritage environments. These challenges motivate the adoption of advanced Digital Twin paradigms that emphasize interoperability, distributed intelligence, privacy preservation, and collaborative learning. Consequently, the following sections examine the application of Digital Twins within cultural heritage contexts and discuss the role of Federated Learning as a key enabler for multi-institutional collaboration.

2.1.3 Application of Digital Twins in Cultural Heritage

The application of DTs in cultural heritage represents a paradigm shift from static digitization approaches to dynamic, data-driven [64]. Traditional heritage documentation methods, such as 3D scanning and photogrammetry, capture detailed representations of physical assets but do not support continuous monitoring or predictive analysis [18, 228]. DTs, in contrast, integrate real-time environmental and operational data to simulate conditions that affect heritage preservation, such as humidity, temperature fluctuations, structural vibrations, and visitor movements. By providing actionable insights, DTs enable heritage managers to anticipate deterioration, optimize conservation interventions, and enhance visitor experiences without compromising the integrity of the assets [19, 30].

2.1.4 Advantages and Challenges in Heritage Context

The integration of DTs in cultural heritage offers numerous advantages. These include real-time monitoring, predictive maintenance, enhanced visitor engagement, improved emergency preparedness, and the ability to simulate operational scenarios without physically interacting with artifacts [19, 207]. However, challenges persist, including technical complexity, high implementation costs, data heterogeneity, and organizational constraints. Ethical and legal considerations related to data ownership, privacy, and intellectual property further complicate the adoption of DTs in heritage contexts [69, 121].

2.1.5 Motivation for Federated Learning in Heritage

Heritage institutions often operate in isolation due to data sensitivity, regulatory constraints, and institutional autonomy. Centralized AI approaches, which require the aggregation of data from multiple sites, are frequently infeasible. Federated Learning addresses this challenge by allowing institutions to collaboratively train machine learning models without sharing raw data [139, 280]. This approach enables multi-institution collaboration while preserving the privacy and sovereignty of heritage data.

In FL, local nodes (e.g., museums) train machine learning models on their private datasets and share only model updates with a central aggregator. The aggregator combines these updates to create a global model, which is then redistributed to the local nodes for further training. This iterative process continues until the model converges [172]. FL not only protects sensitive data but also allows institutions to benefit from collective intelligence, improving model performance and generalizability [155].

2.2 Intelligent and Federated Digital Twins in Cultural Heritage

The digital transformation of cultural heritage management has evolved from simple digitization toward the creation of dynamic, data-driven ecosystems [149]. Digital Twins (DTs), originally developed for industrial applications, are now increasingly being adopted in heritage contexts for real-time monitoring, predictive maintenance, and operational intelligence [30, 241]. These virtual replicas of physical entities provide continuous synchronization between real and digital environments, enabling data-informed decision-making and preservation planning.

As this paradigm matures, the integration of Artificial Intelligence (AI), Internet of Things (IoT), and Federated Learning (FL) frameworks further enhances the intelligence and scalability of digital twins. FL allows multiple institutions to collaboratively train models without exchanging raw data, thereby ensuring data sovereignty and privacy preservation two essential principles in cultural heritage contexts where artifacts, collections, and visitor information are often highly sensitive [139, 280]. When combined, AI and FL enable distributed reasoning and adaptive behavior across interconnected heritage systems, forming the basis for what this dissertation terms the *Federated Digital Twin (FDT)* ecosystem.

2.2.1 Applications of FDTs for Cultural Heritage

The application of FDTs in cultural heritage extends beyond traditional digitization toward active management, conservation, and collaboration [207]. Each local heritage institution can maintain an autonomous digital twin of its assets encompassing structural, environmental, and operational data while contributing to a federated network that learns collectively. For instance, during emergency events such as fires or earthquakes, an FDT framework can support decentralized coordination among multiple museums, enabling real-time decision-making that accounts for spatial configurations, visitor distribution, and environmental conditions [189, 207].

In everyday operations, federated learning enhances predictive maintenance and conservation strategies. For example, distributed models can predict deterioration patterns or microclimate fluctuations by aggregating local insights across multiple sites. Unlike centralized systems, this

approach ensures that each institution maintains full control over its cultural data while benefiting from collective intelligence. The result is an intelligent, collaborative infrastructure capable of supporting sustainable, ethical, and adaptive heritage management. Recent advancements have highlighted the need for real-time coordination in federated digital twin systems, particularly in safety-critical cultural heritage contexts. To address this, Chapter 5 of this dissertation introduces a Real-Time Federated Digital Twin Framework (RT-FDTF), which integrates Worst-Case Execution Time (WCET) profiling, real-time scheduling, and bounded-latency federated synchronization. This framework ensures deterministic decision-making and enhances the operational reliability of distributed digital twin systems in emergency scenarios.

2.2.2 Cognitive Evolution of Digital Twins

While Federated Digital Twins address issues of collaboration and privacy, their intelligence remains largely reactive and dependent on human supervision. The next evolutionary stage *Cognitive Federated Digital Twins (CFDTs)* introduces reasoning, self-adaptation, and ethical decision-making capabilities [154, 293]. Cognitive Digital Twins (CDTs) build upon FDT principles by incorporating cognitive computing mechanisms, reinforcement learning, and semantic understanding, transforming digital twins from static models into autonomous agents capable of contextual learning and continuous improvement [218].

The CFDT model integrates the MAPE-K (Monitor–Analyze–Plan–Execute–Knowledge) loop to enable continuous feedback, situational awareness, and learning across federated nodes. In cultural heritage settings, CFDTs can interpret environmental, structural, and behavioral data to autonomously adapt preservation strategies, optimize energy use, or guide evacuation decisions during emergencies [157]. Importantly, this intelligence is enriched with ethical reasoning inspired by frameworks such as the UNESCO (2022) *Recommendation on the Ethics of Artificial Intelligence*, ensuring that adaptive decisions align with institutional missions and societal values [121, 251].

Through this cognitive extension, digital twins evolve from being passive data systems to proactive, context-aware agents that participate in decision-making processes. This transformation not only enhances operational efficiency but also introduces a human-centered ethical dimension essential for cultural sustainability.

2.3 Digital Twin Maturity and Adaptation in Heritage

Digital Twin maturity models in industry typically assess readiness along dimensions such as connectivity, data integration, computational intelligence, and autonomy [117, 154]. These frameworks provide useful benchmarks for transitioning from descriptive to predictive and intelligent DTs. However, their focus on industrial productivity overlooks socio-technical and ethical dimensions that are fundamental to heritage management [69, 121].

Cultural heritage institutions operate within complex ethical, regulatory, and community-based frameworks. Thus, maturity assessment must include additional dimensions such as data sovereignty, participatory governance, and inclusivity. The proposed framework in this dissertation therefore adapts industrial models to encompass cultural, ethical, and societal indicators, ensuring that technical progress aligns with heritage values and public trust.

In adapting maturity models for heritage contexts, this research identifies four interrelated stages of Digital Twin evolution: *Pre-DT (static models)*, *DT (descriptive and diagnostic)*, *Federated DT (adaptive and collaborative)*, and *Cognitive Federated DT (intelligent and ethical)* as investigated in chapter 3. Each stage corresponds to increasing levels of automation, intelligence, and ethical reasoning.

This progressive model captures the unique trajectory of heritage digital transformation, emphasizing not only technological readiness but also institutional capacity, governance maturity, and cross-sector collaboration. It provides a systematic way to assess how far heritage institutions have progressed toward fully integrated and autonomous digital ecosystems.

2.4 Integration of DT, FL, and AI in Heritage Management

The convergence of DTs, Federated Learning, and Artificial Intelligence provides a unified approach to intelligent and sustainable heritage management. DTs enable real-time monitoring and simulation of physical assets, FL ensures decentralized learning and collaboration, and AI introduces predictive and reasoning capabilities that support decision-making under uncertainty [207, 293]. Together, these technologies enable a resilient ecosystem that supports both preservation and adaptive operations.

In the proposed framework, FL acts as the connective intelligence that allows distributed DTs to exchange learned patterns without exposing raw data. AI models embedded within each twin analyze contextual inputs, such as humidity, structural stress, or visitor flow, to generate adaptive responses. This symbiotic integration facilitates multi-level collaboration across heritage networks local, regional, and global enabling institutions to collectively advance toward intelligent cultural ecosystems.

2.5 Summary of Limitations and Research Opportunities

Despite significant advancements in digital heritage technologies, key challenges persist. Most heritage applications remain static and fragmented, lacking real-time intelligence and interoperability. AI and federated approaches are underexplored, and ethical maturity models tailored for heritage are still in early development. Furthermore, cognitive reasoning a critical enabler of self-adaptive heritage management has yet to be fully realized in current frameworks [30, 121].

Addressing these challenges requires a holistic framework that integrates distributed intelligence, cognitive adaptability, and ethical governance. This dissertation responds to this need by proposing a **Federated and Cognitive Digital Twin Framework** that unites technical scalability, institutional autonomy, and human-centered design. The following chapter introduces the conceptual architecture and operational mechanisms of the proposed framework, demonstrating how it supports both collaborative preservation and emergency response in multi-site heritage environments.

Chapter 3

Systematic Analysis of Digital Twin Applications in Cultural Heritage

Chapter Overview

Building upon the theoretical foundations and technological background presented in Chapters 1 and 2, this chapter provides a comprehensive and systematic analysis of the current state-of-the-art applications of Digital Twin (DT) technology in the domain of cultural heritage. The objective is to consolidate existing research, identify emerging trends, and highlight the gaps and limitations that justify the development of Federated and Cognitive Digital Twin frameworks in the subsequent chapters. The review is structured around a tri-dimensional analytical framework comprising three interrelated perspectives: user-centric applications, enabling technologies, and digital twin maturity levels. This framework facilitates a holistic understanding of how DTs are conceptualized, designed, and implemented to support heritage preservation, management, and public engagement.

Integration of published work this chapter incorporates, with minor editorial modifications, the core content of the article “Digital Twins for Cultural Heritage: A Systematic Analysis of the State of the Art” published by ACM Computing Surveys (CSUR). Figures, tables, and analytical results have been reformatted to comply with the typographic and structural conventions of this dissertation. References have been harmonized with the thesis-wide bibliography, and cross-references have been added to ensure coherence with other chapters. The chapter preserves the integrity and originality of the published work while situating it within the broader objectives of this dissertation.

3.1 Introduction

Digital Twin (*DT*), defined as the interactive connection between a virtual model and its real-world counterpart, operates through an immersive virtual environment that dynamically mirrors the physical world by continuously updating itself based on real-time data [113]. A *DT* serves as a comprehensive representation of an object or system throughout its lifecycle [111], and is increasingly applied across industries for simulation, monitoring, and optimization [239].

A *DT* typically comprises four essential components: (1) the physical entity to be digitalized, (2) its virtual model that mirrors real-world attributes and behaviors, (3) the connection and synchronization mechanisms enabling continuous bidirectional data exchange, and (4) the stakeholders who interact with and benefit from the system [221]. For example, sensors and Internet of Things (IoT)

devices collect real-time data from the physical object and transmit it to the digital model for synchronization, while simulations and insights generated in the virtual space inform decision-making and optimization for the physical entity.

In recent years, *Digital Twins (DTs)* have gained traction in the context of *Tangible Cultural Heritage (TCH)* as powerful tools for preservation, restoration, and visitor engagement [143]. Cultural heritage can be categorized as *intangible* and *tangible*. *Intangible Cultural Heritage (ICH)* refers to practices, languages, knowledge systems, and oral traditions passed through generations [295], while *Tangible Cultural Heritage* encompasses monuments, artifacts, artworks, and culturally significant sites that reflect a community's history and identity [26].

The potential of DTs for cultural heritage aligns with theories of presence and immersion, which suggest that individuals engage with virtual environments in ways similar to their interactions with physical spaces, often experiencing comparable emotional and cognitive responses [226]. These immersive properties make DTs particularly valuable for cultural heritage, where authenticity and user experience are paramount.

Real-world implementations further illustrate the transformative potential of digital twins (DTs) in cultural heritage. For example, UNESCO has promoted DT-based projects to digitally safeguard sites such as the ruins of Palmyra in Syria [65] and the Old Bridge of Mostar in Bosnia-Herzegovina [253], enabling virtual reconstructions after conflict-related destruction. Similarly, the Notre-Dame Cathedral restoration project in France has leveraged DTs created through laser scanning and photogrammetry to guide reconstruction efforts after the 2019 fire [97]. In the museum sector, initiatives such as the Virtual Museums of the European Union (ViMM) [257] and the Smithsonian's 3D Digitization Program [99] demonstrate how DTs enable immersive public access, interactive exhibitions, and long-term conservation planning. These cases highlight the practical relevance of DTs and reinforce the necessity of a systematic review to understand their broader role in tangible cultural heritage.

this thesis presents a *Systematic Literature Review (SLR)* to assess the current state of *DT* applications in *TCH*, focusing on three dimensions: (i) user-centric applications, (ii) enabling technologies, and (iii) maturity levels. This structured approach enables us to (1) evaluate how DTs enhance user engagement and accessibility, (2) analyze enabling technologies such as 3D modeling, XR, IoT, and AI-driven analytics, and (3) assess the maturity of DT implementations, from static digital replicas to adaptive, intelligent systems.

The *objectives* of this review are threefold: (i) to provide a structured overview of DT concepts and their evolution in the TCH domain; (ii) to identify current applications of DTs in cultural heritage, including monitoring, preservation, restoration, and immersive engagement; and (iii) to introduce a tri-dimensional analytical framework that considers user-centric perspectives, enabling technologies, and maturity levels. These objectives define the scope of the review and guide the subsequent analysis.

3.1.1 Motivation and Paper Contributions

The rapid evolution of Digital Twin (*DT*) technology presents significant opportunities for transforming the *Tangible Cultural Heritage (TCH)* domain by enabling dynamic, interactive, and immersive experiences. However, despite increasing interest, research on DT applications in TCH remains fragmented and limited in scope. Most studies address isolated aspects such as 3D re-

construction or visualization, while overlooking comprehensive frameworks that account for user experience, enabling technologies, and maturity levels. This fragmentation hinders the systematic adoption of DTs in cultural heritage practice.

In addition, the literature reveals persistent challenges that justify the need for a systematic review. These include: (i) **technical challenges**, such as real-time synchronization, interoperability of heterogeneous platforms, and scalability; (ii) **legal and ethical challenges**, including intellectual property rights, privacy, and data ownership; (iii) **interoperability challenges**, stemming from the lack of standardized frameworks and protocols for data exchange across institutions; (iv) **organizational challenges**, such as limited institutional capacity and resistance to digital transformation; and (v) **data standardization challenges**, particularly harmonizing formats for 3D models, metadata, and archival datasets. Addressing these challenges requires a holistic analysis that integrates user perspectives, technological enablers, and maturity levels an approach that is currently missing in prior reviews.

The contribution of this thesis can be summarized as follows:

1. Bridge theory and practice: Provide a comprehensive synthesis of how DTs are currently being applied in TCH, moving beyond theoretical definitions to real-world applications in preservation, restoration, monitoring, and immersive engagement.
2. User-centric applications: Examine the extent to which DTs consider diverse stakeholders such as visitors, researchers, and educators focusing on accessibility, inclusivity, and engagement.
3. Enabling technologies: Identify and analyze the technological foundations that support DT development for TCH, including 3D modeling (e.g., photogrammetry, laser scanning), immersive platforms (e.g., *Unity Reflect*, *Unreal Engine*), and data-driven approaches (AI, IoT, XR).
4. Maturity assessment: Propose a structured tri-dimensional framework for assessing DT maturity, ranging from static digital replicas to intelligent, adaptive systems an aspect largely overlooked in previous reviews.
5. Implementation challenges: Systematically categorize the key barriers (technical, legal, interoperability, organizational, and data standardization), offering a roadmap for future research and practical deployment in cultural heritage.

3.1.2 Paper Structure

The paper is structured as follows: Section 3.2 presents the background and related work on *DT* usage in the context of *CH*. In particular, we provide a comparison between our work and other *SLRs* already presented in the literature to show the gaps in the literature highlighting the novelty of our work. Section 3.3 outlines the research methodology, adopted to conduct our *SLR* based on the guidelines reported in [120]. Through a rigorous screening procedure comprising four iterative steps, including a snowball process, we identified and included in our study 108 pertinent papers out of 1271 initially found. These selected papers serve as the foundation for addressing the research questions outlined in Table 3.2.

Sections 3.4, 3.5, and 3.6 respectively analyze the results addressing each research question. In particular, Section 3.4 systematically explores the application of *DT* for *TCH* by integrating technical and user perspectives. The study advances knowledge on how the adoption of *DT* facilitates preservation, education, and stakeholder participation. Section 3.5 examines the enabling technologies of *DT* in *TCH*, focusing on data acquisition technologies, digital twin platforms, and intelligent systems. Subsequently, Section 3.6 examines the maturity level of *DTs* within *TCH* domains.

Section 3.7 discusses the applications of *DTs* from a user perspective, examines enabling technologies, evaluates the maturity levels of *DT* implementations for *TCH*, identifies challenges, and highlights areas where addressing technological and methodological gaps can unlock their full potential as adaptive intelligent systems for cultural preservation and engagement. Section 3.8 explore open challenges and future research directions. Finally, Section 3.9, concludes the paper.

3.2 Background and Related Work

The concept of the *Digital Twin (DT)* had evolved significantly since its introduction in 2002, when Michael Grieves first proposed it in the context of product lifecycle management [87]. Initially, *DT* focused on static virtual representations of physical assets, with limited data integration. Over the years, advancements in IoT, cloud computing, and artificial intelligence transformed *DT* into dynamic, real-time systems capable of continuous data exchange and simulation [148]. By 2010, industries such as aerospace and manufacturing began adopting *DT* for predictive maintenance and operational optimization. For example, Boeing implemented *DTs* to simulate aircraft performance, reducing maintenance costs, while General Electric used *DTs* to optimize turbine efficiency [137]. The 2010s saw widespread adoption across healthcare, energy, and smart cities, leveraging enhanced computational power and machine learning [191]. For instance, Singapore developed a Virtual Singapore *DT* to simulate urban environments and improve city planning by predicting traffic congestion and optimizing public transportation [84]. Between 2020 and 2025, the integration of *DT* with emerging technologies like *extended reality (XR)* and *federated learning* enabled multi-stakeholder applications, including cultural heritage preservation and decentralized systems [10, 109]. The development of *DT* technology has been significantly enhanced through the incorporation of *artificial intelligence (AI)* and *machine learning*, enabling systems to predict and respond to different scenarios in real-time [70]. In sectors such as cultural heritage, this advancement allows for the creation of interactive and immersive digital replicas of museums and historic sites, which improve access and engagement while ensuring the preservation of the physical structures. These applications underscore the powerful impact that *DT* technology can have on tackling intricate challenges across various industries, promoting greater efficiency and resilience [88].

3.2.1 Digital Twin for Cultural Heritage

Digital twin (DT) have become transformative technology in *Tangible Cultural Heritage (TCH)*, enabling the preservation, management, and enhancement of historical sites and artifacts [54]. By creating virtual replicas of physical assets, *DT* facilitates real-time monitoring, predictive maintenance, and interactive experiences for both heritage professionals and the public. Their application in museums and archaeological sites supports data-driven conservation strategies, simulates visitor interactions, and enhances accessibility through virtual tours [236]. These systems integrate

emerging technologies such as IoT, AI, and XR to offer collaborative and immersive solutions that adapt to dynamic environmental and operational challenges [103, 160]. For example, the development of virtual museums leverages DTs to digitize collections, provide customized visitor paths, and optimize exhibition layouts, creating more engaging experiences. This approach is grounded in extensive stakeholder collaboration and lifecycle management, ensuring both the authenticity and sustainability of cultural assets [49]. Recent advancements emphasize scalable, template-based solutions for diverse applications, aligning with broader digitization and heritage preservation [170].

While DTs in general share foundational attributes such as real-time data synchronization, predictive modeling, and simulation capabilities, DTs for TCH demand specialized adaptations tailored to the preservation and engagement with TCH [249]. For instance, DT for TCH emphasizes immersive visualization, historical data integration, and multi-sensory interaction capabilities critical for recreating and preserving cultural narratives but receive less focus in other DT applications [58]. DTs used in industries often prioritize scalability, predictive maintenance, and broad interoperability, which, while relevant, are not central to the mission of DTs for TCH [41]. DT for TCH also requires unique features such as artifact digitization, authenticity preservation, and user-centric designs to enhance accessibility and engagement [175].

3.2.2 Related Work on Digital Twin for Cultural Heritage

Over the past two decades, there has been significant growth in the body of literature addressing the application of DTs to the realm of *tangible cultural heritage* (TCH). The utilization of DT in the management and preservation of TCH assets necessitates the adoption of a collaborative and integrated approach [258]. Currently, a systematic review of the literature has noted a substantial increase in the utilization of DTs within the TCH domain. The following *secondary studies* provide evidence of this trend.

In [101] the authors investigate the theoretical and practical links between DT application in *Heritage Facilities Management* (HFM) from a life cycle management perspective and to signpost the future development directions of DT in HFM. The review results show that DT has been mainly adopted to support decision-making on conservation approach and method selection, performance monitoring and prediction, maintenance strategies design and development, and energy evaluation and management. In their study, authors in [160], focused on digital transformation in leisure activities, especially all kinds of museums, exhibitions, and their DTs. Through a literature search and evaluation, they provide an analytic overview of the necessary concepts and methods involved in a collaborative effort to produce a cost-effective approach to a semi-automatically generated virtual museum. The findings and conclusions of this study concerning the life cycle of a DT essentially relate and are limited to the field of *Heritage 4.0* and virtual museums of all categories. Although these fields offer insights into key aspects of human society, they primarily emphasize the past rather than the future and are significantly less complex compared to the modern challenges and solutions found in *Industry 4.0* subdomains. Authors in [60] studied the digitization efforts in TCH has sparked global interest in their potential implications for preserving, presenting, and managing TCH assets in China. This study summarizes the dimensions and complexity of DT applications on cultural world heritage sites in China, which vary significantly. DT technology has been widely applied to monuments, cultural landscapes, groups of buildings, and sites. The most commonly used methods for collecting graphical and spatial data are 3D laser scanning and photogrammetry, with

their earliest application dating back to the 1980s, coinciding with the inscription of cultural world heritage sites. This thesis investigates the relationship between *DT* technology and heritage conservation, categorizing *DT* techniques into six levels to address three main tasks: data acquisition, visualization, and application. It systematically examines Chinese heritage digitization practices, taking into account variations in complexity and purposes, and highlights typological and chronological characteristics as well as the involvement of stakeholders. The study aims [258] to define the role of *DT* in the conservation of built cultural heritage (BCH). It systematically reviews existing literature to clarify the concept, application areas, and enabling technologies of *DT* in this domain. The study seeks to identify current limitations, such as the lack of standardized frameworks and real-time data integration.

More recently, Mazzetto et al. [54], provided an interdisciplinary synthesis combining *DT* with BIM, IoT, machine learning, and AR/VR for heritage building conservation, identifying interoperability and *DT*–BIM integration as major research gaps. Another systematic review on the management of heritage facilities explores the integration of *DT* and HBIM for conservation and energy management, highlighting structured lifecycle modeling as a key enabler [101].

Recent developments in Museum 4.0 have significantly influenced the adoption of Digital Twin technologies in cultural heritage institutions [149]. Museum 4.0 promotes data-driven, interactive, and user-centered museum environments through the integration of IoT infrastructures, immersive technologies, artificial intelligence, and digital platforms [160]. In this context, Digital Twins support not only digital preservation, but also real-time monitoring, visitor engagement, predictive maintenance, and adaptive exhibition management [79, 160].

Several projects demonstrate the convergence between Museum 4.0 principles and Digital Twin technologies. For example, museums increasingly employ intelligent sensing environments and XR-based interaction systems to personalize visitor experiences while simultaneously collecting operational and environmental data for conservation purposes. These approaches align with the broader transition toward smart and connected museums [80, 101].

In parallel, the CIPA Heritage Documentation framework has played a central role in establishing standardized approaches for cultural heritage digitization [227]. CIPA-based methodologies emphasize accurate geometric documentation using photogrammetry, laser scanning, semantic enrichment, and interoperability standards such as CIDOC CRM. These methods provide reliable digital representations capable of supporting conservation, restoration, monitoring, and long-term archival activities [68, 227].

A representative example is the restoration of Notre-Dame Cathedral following the 2019 fire. The project relied extensively on high-precision laser scanning and photogrammetric datasets that had been acquired before the disaster. These digital resources enabled experts to reconstruct damaged architectural components, assess structural conditions, and support restoration planning with high accuracy [179]. Recent research further expanded this effort through the integration of HBIM, environmental monitoring, and simulation-driven analysis for structural assessment and long-term conservation management [72, 179]. The Notre-Dame project has become a major reference case demonstrating how Digital Twins can support large-scale heritage restoration through the integration of geomatics, computational sciences, and collaborative digital infrastructures.

Meanwhile, our systematic literature review examines the three-dimensional aspects of user-centric applications, enabling technologies, and the maturity level of digital twin utilization in the

Table 3.1: Comparative analysis of prior surveys with respect to the three dimensions considered in this study

Title	Dimension(s) Considered
Digital twin applications on cultural world heritage sites in China: A state-of-the-art overview [60]	None
Digital twin application in heritage facilities management: systematic literature review and future development direction [101]	None
Digital Twins and Enabling Technologies in Museums and Cultural Heritage: An Overview [160]	Enabling technologies
Shaping Digital Twin Concept for Built Cultural Heritage Conservation: A Systematic Literature Review [258]	Enabling technologies
Integrating emerging technologies with digital twins for heritage building conservation: an interdisciplinary approach with expert insights and bibliometric analysis [54]	Enabling technologies and maturity levels
Digital twin application in heritage facilities management: Systematic literature review and future development directions [101]	Enabling technologies

context of *TCH*. This review aims to provide a comprehensive understanding of how these elements interrelate, assess the current state of digital twin technologies in cultural heritage preservation, and identify key factors influencing their effectiveness and future development. Table 3.1 summarizes the objective and key findings of each *SLR* discussed in this subsection. In addition, the Table provides evidence on which of the three dimensions of *DT* in *TCH* analyzed in our study (*user-centric application*, *enabling technologies*, *maturity level*) is addressed by the reported *SLR* allowing to evidence the novelty of our research.

Taken together, the reviewed studies highlight both progress and persistent gaps in the application of Digital Twin (DT) technology to Tangible Cultural Heritage (TCH). Earlier work was focused on static 3D reconstructions and digital documentation, with limited interactivity and real-time data exchange. More recent research, such as DT-based monitoring of church artifacts using IoT sensors, predictive modeling in the Notre Dame restoration project [97], and UNESCO-driven reconstructions of World Heritage sites [253], demonstrates a growing trend toward integrating AI, IoT, and XR for improved preservation and public engagement. However, despite these advances, several gaps remain. Many studies lack user-centered evaluation, particularly in terms of accessibility and adaptive interfaces for diverse audiences. Similarly, the technical integration of AI-driven analytics and IoT-enabled monitoring is often presented in isolation, without a holistic framework for interoperability. Finally, while DT maturity models are increasingly referenced, few works provide structured approaches for assessing progression from static models to intelligent, interactive systems. These limitations collectively underscore the need for our tri-dimensional framework spanning user-centric applications, enabling technologies, and maturity levels, which positions this review as both a synthesis of prior work and a roadmap for future cultural heritage DT development.

Table 3.2: Research Questions

ID	Research Question.
RQ1	How can digital twin technologies enhance the management and experience of tangible cultural heritage?
RQ1.1	How have DT Technologies been adopted in the context of tangible cultural heritage?
RQ1.2	How have digital twins been used to improve stakeholders' experiences?
RQ2	How is digital twin used in the context of cultural heritage?
RQ2.1	What platforms and tools are used for deploying digital twins of cultural heritage for public access?
RQ2.2	What methods are commonly used to capture data for creating <i>digital twins</i> of <i>tangible cultural heritage</i> sites and artifacts?
RQ2.3	How does <i>digital twin</i> intelligence leverage advanced operations for tangible cultural heritage?
RQ3	What is the current maturity level of <i>digital twins</i> in the context of <i>tangible cultural heritage</i> ?
RQ3.1	What is a representative model to assess the maturity of <i>digital twin</i> implementations in <i>tangible cultural heritage</i> ?
RQ3.2	How can the advancement of <i>digital twin</i> implementation be evaluated to assess its maturity in <i>tangible cultural heritage</i> ?

3.3 Research Method

The *Systematic Literature Review (SLR)* presented in this study adheres to the guidelines specifically designed by Kitchenham in [120]. This methodology is composed of three standardised stages, encompassing *planning*, *conducting*, and *reporting*.

3.3.1 Planning the Review

In a *SLR*, the formulation of *Research Questions (RQs)* plays a crucial role in providing a well-defined and organised framework for the review process. Table 3.2 shows the three main research questions we aim to answer with respect to the defined three dimensions: ***user perspective***, ***enabling technologies***, and ***maturity level*** of the *DT* adoption in the context of *TCH*. Research question *RQ1* investigates the application of *DTs* technology in *TCH*, focusing on both technical and *user perspectives* to explore how *DTs* contribute to preservation, education, and user engagement. *RQ2* addresses how *DTs* are used in *TCH*, examining the *enabling technologies*, platforms, and tools for their deployment. Finally, *RQ3* explores the *maturity level* of *DT* in *TCH* by evaluating the model and the current implementation stages and computational intelligence to assess their operational efficiency and effectiveness in preservation and management.

We constructed the search string considering the terms and synonyms presented in Table 3.3. Those terms were extracted by analyzing not only the secondary studies reported in Section 3.2, but also three primary studies we selected as pilot due to their relevance with respect to the defined research questions [166, 199, 278]. We restricted our research to four databases: *Scopus*, *IEEE Xplore*, *ACM Digital Library*, and *Web of Science*. In addition, *Google Scholar* was included to enhance the comprehensiveness of the literature search by capturing a wider range of sources; the

analysis was restricted to the first 100 relevant results to balance breadth with feasibility and to focus on the most pertinent contributions [104]. We tested different strings in the digital libraries to assess their quality based on the relevance of the retrieved papers. For this purpose, we employed two different strategies: i) we checked if the pilot studies were included in the search, ii) each author checked the title and abstract of the first 10 papers in each library's results and discussed their relevance with the other authors. Based on the discussion, we defined the following general search string. The string was then adapted according to the syntax of each library. Details about the search strings we used in each library can be found in our replication package [56]

(((("digital" AND ("twin*" OR "replica" OR "copy")) OR ("virtual" AND ("copy" OR "representation" OR "twin*")))) AND ("cultural heritage*" OR "museum*" OR "gallery*" OR "archeological site*"))

The data extraction plan is an essential component of the *SLR* process as it provides a clear and structured approach to data extraction, which helps to ensure the quality and reliability of the findings. In our *SLR*, we used *Zotero* and *Microsoft Excel* for managing the entire review process composed of four screening stages and reported in Figure 3.1: the *primary screening*, run by the first author of the paper, assessed paper accessibility and language while eliminating duplicates. The *secondary screening*, run by the first and second author of this thesis, evaluated titles, abstracts, introductions, and conclusions. The *tertiary screening*, run by the first author, involved the full reading of the papers for their evaluation for inclusion. Lastly, forward and backward snowballing processes identified relevant studies referenced in included papers and analysed papers citing included [274].

Table 3.4 presents the inclusion and exclusion criteria we adopted in each phase of our study. Note that we restricted our research from 2002 (*P-IC3*), as *Grieves* defined the concept of the digital twin as a virtual instance of physical assets capable of continuously mirroring them [87]. During the *secondary screening*, we discovered papers with titles in English, but the rest of the content was in another language. Thus, we applied the same criteria as in the primary screening (*P-IC1*, *S-IC3*).

The inclusion and exclusion criteria were used during the *secondary* and *tertiary* screening to decide the papers to be moved to the next step while assessing their quality. On one hand, in the secondary screening, we used the exclusion criteria to exclude any paper for which at least one exclusion criteria was satisfied. For example, papers discussing about *digital twin* (*S-EC2* not satisfied) but consisting of just an abstract (*S-EC1* satisfied) have been excluded. On the other hand, the inclusion criteria were used to qualitatively assess the papers under analysis. In this case, we assigned a value (0, 0.5, 1) to each inclusion criteria based on how much the paper satisfies the inclusion criteria under consideration. Thus, we moved to the tertiary screening all the papers

Table 3.3: Search Keywords: Terms and Synonyms

Term	Synonyms
Digital Twins	digital twin, digital replica, digital copy, virtual copy, virtual representation, virtual twin
Cultural Heritage	cultural heritages, museum, museums, galleries, gallery, archeological site, archeological sites

Table 3.4: Inclusion and exclusion criteria used

Screening phase	Inclusion criteria	Exclusion criteria
Primary Screening	P-IC1: Paper is written in English (based on titles). P-IC2: Paper is not a duplicate from multiple sources. P-IC3: Year of publication from 1/1/2002. P-IC4: Full text is accessible.	P-EC1: The paper is not written in English. P-EC2: Full text not accessible. P-EC3: Duplicated papers.
Secondary screening	S-IC1: The paper address <i>digital twin</i> in <i>cultural heritage</i>. S-IC2: The paper conceptualizes <i>digital twin</i> maturity level S-IC3: The full paper is written in English. S-IC4: Studies directly related to the research questions.	S-EC1: The content of the paper is only an abstract. S-EC2: The paper mention <i>digital twin</i>, but it cannot help to answer the RQs.
Tertiary screening	T-IC1: Addressing <i>digital twin</i> in <i>cultural heritage</i> and museums. T-IC2: The paper conceptualize <i>digital twin</i> maturity level. T-IC3: The paper relates to the research questions.	T-EC1: None of the inclusion criteria is satisfied.
Backward screening	Tertiary screening.	
Forward screening	Tertiary screening.	

having a score greater or equal to 2.5. The same scoring process has been adopted in the tertiary screening. In particular, we excluded all the papers not satisfying at least one inclusion criteria. We moved to the full reading phase all the papers scoring more than 1. Despite this rule, the minimum score registered for the paper included in the tertiary screening was 2. Indeed, while it was easy to evaluate if the papers address *digital twin* in *cultural heritages* (*T-IC1*) and if the paper relates to the research questions (*T-IC3*), more difficult was to understand if the paper conceptualize the maturity level of *digital twins* in the context of *cultural heritage*. More information about each phase as well as the quality metrics used for papers eligibility can be found in our replication package [56].

**Figure 3.1:** Screening Phase Description and Numbers

Table 3.5: Selected Parameters for the Study

ID	Name	ID	Name
RQ1.1.P1	Types of artworks [174]	RQ2.3.P1	Activity level [58, 195]
RQ1.1.P2	Purpose of the tangible cultural heritage digital twin [52, 166]	RQ2.3.P2	Functional layer of DT [58]
RQ1.1.P3	Digital twin techniques [258]	RQ3.1.P1	Model update frequency [44]
RQ1.2.P1	Immersive technologies [216]	RQ3.1.P2	Twin models [58]
RQ1.2.P2	Digital heritage user experience [125]	RQ3.1.P3	Model security [142]
RQ1.2.P3	End users of digital heritage [29]	RQ3.1.P4	Verification and Validation of DT Models [259]
RQ2.1.P1	3D modeling software [102]	RQ3.2.P1	Stage of DT implementation [58, 132]
RQ2.1.P2	Digital twin platforms [209]	RQ3.2.P2	Connectivity mode [58, 132]
RQ2.2.P1	Data-acquisition technology [29]	RQ3.2.P3	Computational Intelligence [44]
RQ2.2.P2	Types of data [29]		

3.3.2 Conducting and Reporting the Review

According to the Kitchenham et al. [120] guidelines, the conduction phase involves extracting relevant data from the selected studies and synthesizing the findings to address the *RQs*. Figure 3.1 summarizes the search results for each digital library, indicating the number of papers retrieved from each library. In addition, Figure 3.1 presents the numerical outcome after performing each screening phase. Finally, we established 20 *RQ*-related parameters, as shown in Table 3.5. We inferred those parameters and related attributes from the literature *on DTs* [58, 88], and are intended to capture all necessary data for addressing the *RQs* and their sub-questions. The descriptions of the selected parameters and attributes are reported in our replication package [56].

Data synthesis in a systematic literature review (SLR) aims to analyse and summarise the findings from the primary studies included in the review. This phase entails extracting data from the studies, organising and categorising it, and synthesising the information to address the research questions (RQs). In our *SLR*, we integrated the thematic analysis method [245] for qualitative data synthesis and used meta-analysis [53] for quantitative data synthesis. The following sections of this thesis present the analysis results to address the given *RQs*.

3.4 Digital Twin Technologies for Tangible Cultural Heritage Management (RQ1)

The main objective of this section is to comprehensively address the application of digital twin technology to tangible cultural heritage by integrating both technical and user perspectives. The investigation will contribute to understanding how digital twins can serve as tools for preservation, education, and enhanced user engagement while ensuring accessibility and stakeholder participation, thus addressing the objectives of RQ1.1 and RQ1.2. This aligns with the main research question RQ1: “How can digital twin technologies enhance the management and experience of tangible cultural heritage?”

3.4.1 Adoption of Digital Twin Technologies (RQ1.1)

This section aims to answer RQ1.1, “How have DT Technologies been adopted in the context of tangible cultural heritage?”. The objective of this research question is to explore the rationale behind the adoption of *Digital Twin (DT)* technologies for *Tangible Cultural Heritage (TCH)* and to investigate the types of cultural artifacts where digital twins are applied to enhance services. To address this research query, we’ve utilized specific elements. These include **types of artwork**, **DT techniques**, and **purpose of the TCH digital twin**. These components are crucial in the adoption of *DTs* within the *TCH* realm. They collectively enable the generation of immersive and educational renditions aimed at conserving, safeguarding, and showcasing our diverse cultural legacy for posterity. In the realm of *TCH* studies, the digital preservation of artworks emerges as a pivotal undertaking. This domain focuses on implementing strategies geared towards safeguarding and proliferating our *TCH*, with the overarching goal of ensuring its accessibility and enduring significance for future generations [58].

Diverse **types of artworks**, such as *buildings*, *paintings*, *drawings*, *monuments*, and *sculptures*, are under consideration for digitalization. The complete list of possible values associated with the *types of artworks* parameter can be found in our replication package [56]. Findings from our study reveal that among the selected 108 studies, buildings emerge as the most digitized form of artwork, with 61 ($\approx 57\%$) studies focusing on them. *Sculptures* closely follows with 33 ($\approx 30\%$) studies, while *monuments* account for 11 ($\approx 10\%$) studies, and *paintings* and *drawings* are represented by 2 ($\approx 2\%$) and 1 ($\approx 1\%$) studies, respectively. This prioritization is influenced by the widespread presence of *buildings* and *sculptures* in cultural institutions worldwide and their capacity to represent various societal aspects. *Monuments*, *paintings*, and *drawings* receive comparatively less focus, potentially due to their varied sizes and mediums [58]. This prioritization highlights the importance of preserving architectural and sculptural heritage, which embodies diverse cultural narratives. This choice reflects the commitment to digitizing these forms of art, ensuring that they remain accessible to future generations, and safeguarding their significance in *TCH* preservation [167].

The **purpose of DT for TCH** is to enable the *preservation*, *conservation*, and *restoration* of cultural artifacts through virtual replication of cultural artifacts and artworks. This technological advancement allows us to recreate and safeguard *TCH* in a digital format. In this context, our study examines the rationale behind digital twins, emphasizing their importance in *preservation*, *conservation*, and *restoration* efforts for *TCH* to play the crucial roles in safeguarding and maintaining the diversity of human history, art, and traditions [273]. An analysis of included study papers reveals that *preservation* is the primary purpose, with 63 ($\approx 58\%$) studies dedicated to this aspect, followed by *conservation* 29 ($\approx 27\%$) and *restoration* 15 ($\approx 14\%$), while 1 ($\approx 1\%$) study does not explicitly state its purpose of digital twin utilization. Our study highlights *preservation* as the primary driving force behind the creation of virtual replicas of museums and cultural artifacts. It emerges as the most prominent motivation, given its central focus on preventing deterioration and preserving the original condition of these invaluable treasures and environments. By meticulously documenting every detail in digital form, *preservation* ensures that the essence and integrity of these artifacts remain intact for future generations to experience and appreciate [62]. In contrast, *conservation* steps in to address existing damage, employing strategies to manage and mitigate further degradation [142]. When creating virtual replicas, *preservation* ensures meticulous documentation of the current state, capturing every detail with precision. Meanwhile, *conserva-*

Table 3.6: Co-occurrence of studies on types of artworks (RQ1.1.P1) and Purpose of the CH digital twin (RQ1.1.P2).

	Preservation	Conservation	Restoration
Building	[9, 40, 41, 47, 72, 76, 91, 110, 114, 116, 123, 124, 126, 133, 158, 159, 165, 181, 185, 193, 196, 200, 215–217, 229, 230, 244, 248, 264, 276, 279, 281, 289, 291]	[1, 59, 93, 100, 108, 128, 130, 132, 150, 167, 256, 265, 278, 281, 285, 287, 296]	[24, 31, 43, 52, 96, 98, 147, 179, 197]
Sculpture	[7, 11, 28, 35, 39, 49, 55, 73, 89, 135, 138, 140, 166, 178, 182, 226, 255, 260, 269]	[74, 80, 118, 161, 188, 199, 237, 263]	[171, 224, 282]
Painting	[184]	[153]	N/A
Drawing	N/A	[81]	N/A
Monuments	[95, 198, 220, 247, 262, 284]	[16, 22]	[180, 210, 270]

tion efforts focus on digital *restoration*, meticulously repairing any flaws or imperfections present in the scanned artifacts. Restoration, a facet closely intertwined with *conservation*, involves digitally recreating missing elements to restore the artifact to its original splendor. Together, *preservation*, *conservation*, and *restoration* converge in the digital realm, where advanced technologies enable the faithful *preservation* and *restoration* of *TCH* for future generations to appreciate and cherish [263].

In the digitalization of *TCH*, a diverse array of **DT techniques** have emerged, offering innovative ways to capture and preserve the essence of historical sites, artifacts, and cultural treasures. These techniques can be broadly categorized into *data-driven models* that focus on learning from data, *physics-based models* that simulate physical principles, and *hybrid models* that combine both approaches for a more robust representation within the *DT* framework [167]. A recent study sheds light on the evolving landscape of *DT* techniques within the *TCH* sector, revealing a predominant utilization of data-driven approaches over the last two decades because *data-driven DTs* offer greater adaptability to changes in the cultural heritage environment. As new data becomes available, or heritage operations evolve, the *DT* can be updated and refined to reflect these changes, while *physics-based DTs* often require complex mathematical models and extensive computational resources to simulate physical phenomena accurately. This flexibility allows for ongoing optimization and improvement of user experiences and operational efficiency [55, 261]. This trend is confirmed by our study, revealing a notable prevalence of *data-driven DT techniques*, with 74 ($\approx 68\%$) out of 108 selected studies featuring their utilization. In contrast, *physics-based techniques* were employed in 5 ($\approx 5\%$) out of the included studies, and *hybrid DT techniques* were utilized in 20 ($\approx 19\%$) of the studies. Notably, 9 ($\approx 8\%$) studies within the pool did not explicitly specify the technique employed, indicating the evolving and sometimes complex nature of *DT* applications within the *TCH* sector. The decision between employing *data-driven* or *physics-based* approaches should be guided by the particular research goals, the accessibility of data, and the characteristics of the *TCH* elements being studied [204].

Answer to Research Question 1.1 The pattern, derived from the information presented in Table 3.6, underscored the prevalence of digitalization practices geared towards the preservation of *TCH*. This trend suggests a concerted effort within the *TCH* sector to leverage *DT* technologies to safeguard and immortalize *artworks* such as sculptures, buildings, monuments, paintings, and drawings. Among the observed patterns, a significant correlation was found between the digitalization of buildings and sculptures and the practices of preservation, conservation, and restoration, with the latter occurring in descending order. Conversely, paintings, drawings, and monuments demonstrated weaker associations with these preservation, conservation, and restoration methods. Moreover, the robust correlation between preservation objectives and the digitalization of artworks, particularly buildings, and sculptures, underscores the overarching theme of leveraging *DT* technologies for *TCH* conservation. Here, preservation encompasses not only physical artifact protection but also archival and digital conservation aspects. This integration of preservation and digitalization signifies a comprehensive approach to ensuring the endurance and accessibility of *TCH* assets. Furthermore, the observed co-occurrence patterns indicate that the adoption of *DT* technologies in the *TCH* domain is largely influenced by the nature of the artwork and the intended preservation objective. Buildings and sculptures dominate the adoption landscape because they typically require continuous monitoring, detailed structural documentation, and long-term conservation planning. These artworks also benefit significantly from data-rich digital environments that support visualization, simulation, and condition assessment. In contrast, paintings and drawings exhibit comparatively lower levels of *DT* adoption due to the complexity associated with capturing fine-grained material characteristics and the limited integration of real-time monitoring mechanisms. Additionally, the strong prevalence of data-driven approaches demonstrates the growing reliance on sensor data, image-based reconstruction, machine learning, and 3D modeling techniques to create scalable and adaptive *DT* solutions for cultural heritage applications. This trend reflects the increasing shift toward intelligent and responsive heritage management systems capable of supporting sustainable preservation and enhanced cultural accessibility.

The investigation conducted in this section delved into the exploration of parameters that exhibited more frequent occurrences and significant patterns within our research findings. The observed patterns were then organized into, revealing noteworthy insights. Among the patterns identified in Table 3.7, digitalization for preservation purposes and the application of data-driven *DT* techniques emerged as the most recurrent themes in the context of our study. The pattern, derived from the information presented in Table 3.7, underscored the prevalence of digitalization practices geared toward the preservation of *TCH*. Data-driven techniques often require less computational resources compared to physics-based models, making them more scalable for large-scale applications like visitor flow analysis in *TCH* sites [216, 226].

3.4.2 Immersive Technologies for Improving Stakeholder Experience (RQ1.2).

This subsection aims to answer RQ1.2, “*How digital twins have been used for improving stakeholders experience?*”. To assess the effectiveness of digital twin applications in overcoming physical limitations, expanding audience reach, and fostering deeper connections with *TCH* in alignment with the main research question objectives, we consider parameters such as immersive technologies, access devices, and end users. These elements facilitate immersive and interactive engagement with cultural artifacts and exhibits, significantly enhancing stakeholders’ experiences within cultural

Table 3.7: Co-occurrence of Purpose of the TCH digital twin (RQ1.1.P2) and digital twin techniques (RQ1.1.P3).

	Preservation	Conservation	Restoration
Data-driven DT	[11, 28, 35, 39, 40, 49, 55, 72, 73, 76, 89, 90, 110, 116, 122–124, 135, 138, 140, 158, 159, 165, 178, 181, 182, 184, 185, 193, 200, 215–217, 220, 226, 229, 244, 248, 255, 260, 262, 264, 269, 281, 284]	[1, 16, 59, 74, 80, 81, 100, 118, 130, 150, 153, 161, 167, 188, 237, 254, 256, 285, 296]	[24, 31, 43, 96, 98, 171, 179, 180, 224, 270]
Physics-based DT	[41]	[132, 199]	[197]
Hybrid DT	[7, 47, 91, 95, 126, 166, 198, 230, 247, 276, 279, 289]	[93, 263, 265, 278, 287]	[147, 210]

heritage sites.

In the pursuit of improving stakeholders experiences, in the context of *DTs*, researchers, and practitioners have explored **immersive technologies** such as *Virtual Reality (VR)* and *Augmented Reality (AR)*, *images*, and *videos* [216]. These *techniques* aim to captivate and engage users, offering deeper insights and fostering emotional connections. The toolbox of techniques ranges from the use of *images* to cutting-edge technologies like *VR* and *AR*. This study investigates the prevalence and effectiveness of these methods in enhancing user experiences. Among the 108 studies examined, three primary methods emerged: *images*, *VR* and *AR*, and *videos*. The most widely used type, with 59 ($\approx 55\%$) out of included studies, were *images*. It serves as a visual aid, helping convey information, set the ambiance, and evoke emotions. On the other side, *VR* and *AR* are employed in 39 ($\approx 36\%$) studies respectively. These technologies take stakeholders on immersive journeys, allowing them to explore virtual realms, interact with digital replicas, and experience historical or artistic contexts in novel ways. Recent frameworks, such as those proposed by Anwar et al.[10], emphasize the integration of XR and metaverse technologies to enhance emotional engagement, motivation, and learning outcomes in cultural heritage education. These models advocate for immersive storytelling, gamified environments, and social interaction features such as avatars and collaborative exploration that deepen user experience and accessibility. Moreover, standardized UX evaluation methodologies like ISO 9241-210 are increasingly applied to assess the effectiveness of immersive heritage platforms. While *videos* are a less common choice among the techniques explored in this study, their capacity to add motion, storytelling, and auditory elements to the user experience is not to be underestimated. The limited use of *videos* in 2 ($\approx 2\%$) out of the total included studies suggests that there is potential for further exploration of this medium to enhance stakeholders' engagement. The remaining 8 ($\approx 7\%$) studies did not explicitly specify the methods or techniques employed to enhance stakeholders' experiences, highlighting an area for further research and innovation. The research findings indicate that *images* have been the predominant and extensively utilized method to enhance the experiences of stakeholders in recent decades. This prevalence highlights the lasting

significance of visual content in conveying information, establishing the atmosphere, and evoking emotional reactions from stakeholders. Furthermore, the relatively lower prevalence of *VR* and *AR* compared to *images*, in our study, suggests a burgeoning interest in harnessing these technologies to craft increasingly immersive and interactive user experiences.

The utilization of *DT* technology in the digitization of heritage materials emphasizes the critical importance of the **devices** employed to enhance the user experience [167]. The quality of digital interactions is substantially influenced by individual preferences and the devices utilized as users engage with these digitized assets [288]. This research explores the complexities of user experience within the digital realm, examining the impact of *computers*, *smartphones*, and *tablets*. A notable aspect of this investigation is the preference among users for *computers*, primarily driven by the ubiquitous utilization of web-based platforms. This trend is underscored by the inclusion of 62 ($\approx 58\%$) studies out of a total of included studies, with 37 ($\approx 34\%$) studies focusing on smartphone accessibility and 9 ($\approx 8\%$) on tablet usage. In this study, we find that users exhibit a noticeable preference for *computers* as their primary means of accessing content. This preference is driven by several factors, one of which is the widespread use of web-based platforms that are often optimized for computer interfaces. Furthermore, the prioritization of experimental approaches within this domain signifies a collective effort to enhance user experiences. Researchers and practitioners are committed to gaining valuable insights into user preferences and behaviors, intending to achieve platform-independent outcomes. The dedication to enhancing user interactions with historical, cultural, and artistic artifacts serves as a testament to the importance of preserving and sharing our rich heritage in the digital age [89].

This study also explores the dynamic landscape of end-users within the digital heritage domain, with a focus on the application of *DTs*. It explores the various categories of **end-users**, ranging from the *public* to *researchers*, *students*, and *educators*. The digitalized *TCH*, enriched by *DT* technology, offers a valuable resource that caters to a diverse audience. This investigation highlights the significance of digitized *TCH* as a versatile and accessible resource. It provides engagement opportunities for the *public*, research possibilities for *scholars* and *researchers*, and educational experiences for *students* and *educators*. The analysis of the data reveals that there is a broad focus in the studies, with 69 ($\approx 64\%$) primarily catering to the interests and needs of the *general public*. Additionally, 33 ($\approx 30\%$) studies accommodate the requirements of *researchers*. However, the emphasis on *students and educators* is relatively limited, with only 6 ($\approx 6\%$) study specifically addressing their unique needs out of the total included studies. As a result, it becomes evident that recent research on *DTs* within the *TCH* domain lacks sufficient emphasis on distinct user groups. While the majority of the studies are dedicated to catering to the general public and researchers, there is a notable gap in addressing the needs of students and educators. This area presents an opportunity for future research and development to enhance educational experiences and make the wealth of digitalized *TCH* more accessible and impactful for educational purposes.

Answer to Research Question 1.2 This part of the study proposed to enhance the experience of stakeholders in the digital heritage by identifying key parameters, including *immersive technologies*, *end-users of digital heritage*, and *access devices*. This aimed to explore user engagement dynamics and provide insights into the evolving landscape of digital interactions with cultural artifacts. From Table 3.8 is evident that the general public constitutes a major audience for digitized *TCH*, and

Table 3.8: Co-occurrence of studies on access devices (*RQ1.2.P2*) and end users of digital heritage (*RQ1.2.P3*).

	General public	Researchers	Students and educators
Computers	[7, 11, 24, 49, 52, 55, 72, 81, 89, 90, 96, 108, 110, 118, 123, 132, 159, 161, 167, 178, 180, 181, 197, 217, 226, 247, 248, 262, 269, 270, 278, 285, 288]	[1, 40, 41, 59, 73, 116, 124, 130, 133, 150, 158, 166, 179, 185, 193, 196, 199, 200, 230, 237, 244, 256, 263, 265, 279, 281]	[39, 74, 100, 188]
Smartphones	[16, 28, 31, 35, 47, 80, 91, 93, 95, 98, 114, 126, 128, 135, 138, 147, 171, 182, 184, 198, 215, 216, 220, 229, 255, 260, 276, 282, 284, 289]	[140, 254, 287, 291]	[43, 153]
Tablets	[22, 76, 122, 165, 210, 224]	[9, 264, 296]	N/A

using computers as their primary tools, individuals from the public domain engage with images, indicating a widespread embrace of digital platforms for cultural exploration. A noteworthy observation from our research is the enduring prevalence of images as the primary content accessed by users in the digital realm. Despite the advancements in virtual technologies, traditional forms of visual representation remain the dominant means through which the public accesses and appreciates digitized cultural artifacts. This underscores the enduring power of visual narratives in conveying the essence of *TCH*, even in a digital context.

Immersive technologies, facilitated through computer-based interactions, emerged as a significant channel for the public to explore digitized *TCH* [180]. The utilization of *VR* and *AR* in the context of immersive technology was particularly notable. Our findings indicate a close affinity for (*VR*) and (*AR*) technologies, highlighting their role in enhancing the immersive experience for the public, as depicted in Table 3.9. The integration of these technologies adds a layer of interactivity and realism, providing users with a more engaging and dynamic exploration of cultural artifacts.

3.4.3 Requirements of the Multidisciplinary User Community

Building on the empirical analysis presented in the previous subsection, which highlighted how digital twin applications in cultural heritage are accessed through diverse immersive technologies, devices, and end-user groups, it becomes essential to translate these observed interaction patterns into concrete system requirements. In particular, understanding how different stakeholders engage with digital cultural heritage provides a foundation for defining the functional and non-functional requirements of cultural heritage digital twin ecosystems.

Digital twin ecosystems for cultural heritage must therefore address the heterogeneous requirements of multidisciplinary stakeholders who interact with heritage assets from different operational, scientific, educational, and cultural perspectives [17].

Researchers and conservators typically require access to high-resolution documentation, tempo-

Table 3.9: Co-occurrence of studies on Access devices (*RQ1.2.P2*) and immersive technologies (*RQ1.2.P1*).

	General public	Researchers	Students and educators
VR/AR	[11, 24, 31, 35, 47, 49, 55, 80, 91, 95, 96, 110, 114, 122, 126, 128, 138, 147, 159, 171, 180, 181, 197, 198, 215, 216, 248, 270, 276, 278, 282, 284, 289]	[1, 40, 150, 199, 263, 291]	N/A
Images	[7, 16, 22, 28, 72, 76, 81, 90, 93, 98, 108, 118, 123, 132, 135, 161, 165, 167, 178, 182, 184, 210, 217, 220, 224, 229, 247, 255, 260, 269, 285]	[9, 73, 90, 116, 124, 133, 140, 179, 185, 193, 200, 230, 237, 244, 254, 256, 264, 265, 279, 281, 287, 296]	[39, 43, 100, 153, 188]
Videos	[89]	[59]	N/A

ral monitoring data, restoration histories, and analytical tools capable of supporting conservation planning and scientific investigation. Archaeologists and curators additionally require contextual interpretation mechanisms capable of associating artifacts with provenance, excavation records, semantic relationships, and historical narratives [103].

Site managers and policymakers often prioritize decision-support functionalities, including sustainability indicators, risk assessment dashboards, maintenance forecasting, and regulatory reporting mechanisms. In contrast, educators, visitors, and local communities generally expect intuitive interfaces, immersive storytelling experiences, multilingual accessibility, and meaningful engagement with both tangible and intangible cultural narratives [18].

These heterogeneous requirements demonstrate that cultural heritage digital twins should not be interpreted exclusively as technical replicas of physical assets. Instead, they should function as intelligent and user-oriented knowledge environments capable of integrating semantic heritage representations, collaborative annotation mechanisms, personalized visualization modalities, provenance-aware information access, and participatory knowledge contribution workflows [192].

3.5 Technology Usage of Digital Twin in Tangible Cultural Heritage (RQ2)

This section addresses three sub-research questions focused on data acquisition technologies, digital twin platforms, and digital twin intelligence. These elements collectively answer the main research question of RQ2: “*How is digital twin used in the context of tangible cultural heritage?*” The objective is to evaluate the current use of digital twin technologies and their enabling technologies in the preservation and management of tangible cultural heritage.

3.5.1 Tools and Platforms used for the Deployment of Digital Twin (RQ2.1)

This subsection seeks to address RQ2.1, “*What platforms and tools are used for deploying digital twins of tangible cultural heritage for public access?*”. This sub-RQ aims to support the objectives

of the main research question by investigating key tools and platforms to understand the current utilization of *DT* in *TCH*. To answer this research question, we have employed two crucial parameters: **3D modeling software** and **digital twin platforms**. The selection of these parameters is based on their fundamental roles in the creation and application of *DT* for *TCH*. **3D modeling software** is essential for generating accurate and detailed 3D models of *TCH* sites, which form the basis for virtual reconstructions. **Digital twin platforms** are critical for developing immersive *DT* environments, as they enable dynamic simulations and real-time interactions, enhancing user engagement with historical reconstructions[209]. Recent studies, such as Anwar et al.[10], propose a structured architecture model for XR and metaverse applications in cultural heritage. This model guides the development of immersive educational environments by integrating real-time simulation, user interaction layers, and metadata-driven content delivery. Incorporating such frameworks into *DT* platforms can enhance scalability, interoperability, and user engagement. In the expansive domain of *DT* creation and advancement, a plethora of **3D modeling software** options are available, each offering unique features and capabilities. Prominent among these are *Autodesk*, *Blender*, *SketchUp*, *Reality Capture*, and *3DGIS tools* [102]. Delving into the specifics, our research reveals that *Blender* emerges as the front-runner among the utilized 3D modeling tools, being the choice in 21 ($\approx 19\%$) of the studies. Following closely is *Autodesk*, making its mark in 20 ($\approx 19\%$) studies, while *SketchUp* and *3DGIS* each contribute to 2 ($\approx 2\%$) studies. *Reality Capture*, on the other hand, features in 8($\approx 7\%$) of the total studies. This distribution sheds light on the current landscape, indicating that despite the array of available 3D modeling tools, the practical application in *DT* research remains somewhat limited. This subtle exploration leads to a significant inference that the predominant focus in *DT* research appears to be concentrated on the theoretical or conceptual level, with a notable gap in the practical application of 3D modeling tools to create *DTs*.

The intricate landscape of *DT* implementation, and the utilization of specialized **platforms** like *Unity Reflect*, *Unreal Engine* stand out to materialize *DTs*, transforming them into virtual replicas of physical objects [209]. However, our comprehensive study unveils a striking observation: only a modest fraction, 29 ($\approx 27\%$) studies out of the total included studies, incorporate these platforms in their *DT* endeavors. Delving into the specifics, *Unity Reflect* emerges as the predominant choice, featuring in 24 ($\approx 22\%$) of the 108 total included studies. Following in the hierarchy is *Unreal Engine*, making its mark in 4($\approx 4\%$) studies, while *3DGIS* covers a single ($\approx 1\%$) study. This distribution illuminates a prevailing scenario where the practical application of *DT* platforms remains somewhat limited, particularly within the *TCH* domain. This observation prompts a significant inference: the deployment of *DTs*, as facilitated by dedicated platforms, is still in the early stages of practical application. The limited representation of studies incorporating these platforms indicates a gap between the theoretical understanding of *DTs* and their tangible, hands-on use. Looking ahead, the future holds promising prospects for a more extensive exploration of *DT* platforms in the context of *TCH*. The study serves as a foundational step, shedding light on the current limitations and providing a springboard for future investigations. The envisioned exploration encompasses a more in-depth understanding of how platforms like *Unity Reflect* and *Unreal Engine* can be instrumental in the preservation, conservation, and restoration efforts within the *TCH* domain.

Answer to Research Question 2.1 In *DTs*, the initial steps involve choosing *3D modeling tools* and *DT platforms* to create virtual replicas of physical objects, especially in *TCH* contexts. These

tools hold immense importance, particularly in the development of virtual heritage sites, where they serve as essential components. The diverse landscape of *DT* creation offers numerous options for 3D modeling software alongside implementing platforms, each offering distinct features and capabilities. However, the study underscores that the practical application of *DT* platforms remains somewhat limited, particularly within the *TCH* domain. This observation leads to a significant inference: the deployment of *DTs*, facilitated by dedicated platforms, is still in its early stages of practical application. The limited representation of studies integrating these platforms indicates a gap between the theoretical understanding of *DTs* and their tangible, hands-on implementation. Considering these findings, we recommend the adoption of advanced platforms such as Unity Reflect and Unreal Engine with advanced 3D modeling tools. Furthermore, if feasible, we suggest extending this recommendation to include cloud-based platforms such as *AWS* and *Azure DT*. The inclusion of cloud-based platforms offers several advantages in terms of scalability, flexibility, and accessibility. These platforms provide powerful computing resources that enable the processing of large datasets and real-time simulations, which are essential for managing complex digital twin environments.

3.5.2 Digital Twin Data Acquisition Techniques (RQ2.2)

This subsection is dedicated to exploring RQ2.2, “*What methods are commonly used to capture data for creating digital twins of tangible cultural heritage sites and artifacts?*” This sub-RQ focus is on examining the data acquisition techniques employed in the creation of *DT*, along with the types of data essential for accurately representing *TCH*.

This section highlights the techniques for **data acquisition** from physical artifacts and transforms them into virtual representations [263]. Our study explores the efficacy of well-known data acquisition techniques such as *photogrammetry*, *3D laser scanning*, and *sensor-based* techniques. Our analysis reveals that *Photogrammetry* predominates, with 44 ($\approx 41\%$) out of the total included studies. *3D laser scanning* follows closely, with 34 ($\approx 31\%$) studies employing it. In contrast, data collection based on *sensors* is slow, with only 30 ($\approx 28\%$) studies using this approach. These statistics underscore the pivotal role of *3D laser scanning*, renowned for its precision and efficiency in capturing 3D data. Additionally, *photogrammetry* shines as a widely adopted technique in the field. However, *sensor-based* data collection is less emphasized, as indicated by our study’s findings. The statistical distribution of these data acquisition techniques within the selected studies sheds light on their relative significance. *3D laser scanning* and *photogrammetry* are the primary choices for acquiring data from physical artifacts and converting them into digital forms. Their efficiency and precision make them invaluable for capturing the intricacies of *TCH* items. However, the relatively limited use of *sensor-based* data collection methods indicates potential for further development and application in this domain. *Sensor-based* techniques offer unique advantages, including non-contact scanning and real-time data acquisition, which can be leveraged for specific applications within the *TCH* digitalization landscape [41].

In the field of developing *DTs* for *TCH* digitalization, the acquisition of diverse **data types** is essential for accurate representation. These types encompass *visual data*, *technical data*, and *metadata* [21]. Analysis of the included studies reveals that a substantial majority, 62 ($\approx 58\%$) out of 108, primarily rely on *visual data*. The preference for *visual data* is often attributed to its simplicity in comparison to other data types. While *visual data* is undeniably valuable, fostering advanced and thorough research necessitates prioritizing the utilization of *metadata* and *technical data*

Table 3.10: Co-occurrence of studies on data-acquisition technology (*RQ2.2.P1*) and Type of data (*RQ2.2.P2*).

	Photogrammetry	3D laser Scanner	Sensor
Visual data	[11, 47, 49, 55, 80, 81, 89, 91, 93, 108, 110, 114, 118, 122, 128, 133, 147, 167, 178, 182, 197, 198, 224, 226, 229, 244, 247, 260, 270, 276, 284, 289]	[22, 24, 31, 35, 52, 59, 90, 123, 126, 135, 150, 159, 165, 171, 180, 184, 193, 215–217, 248, 262, 269, 279, 282, 285]	[7, 16, 95, 278]
Metadata	[28, 153, 185, 296]	[1, 132, 158, 199]	[9, 39, 41, 74, 98, 100, 130, 140, 188, 200, 254, 256, 263, 265]
Technical data	[73, 76, 138, 161, 166, 181, 255]	[43, 72, 96, 124, 196, 210, 287, 288]	[40, 116, 179, 230, 237, 264, 281, 291]

types [130]. Our findings indicate that the selected studies related to *DTs* in the *TCH* context do not adequately emphasize the inclusion of *metadata* and *technical data* types, with only 23($\approx 21\%$) studies utilizing *technical data* and *metadata*, which is also 23($\approx 21\%$). The study’s findings reveal a gap in the emphasis placed on *technical data* and *metadata* within the selected studies. *Technical data* encompasses information related to the artifact’s materials, dimensions, and structural properties [196]. *Metadata* includes contextual information, such as historical significance, origin, and provenance [256]. These data types provide essential insights that deepen our understanding of the artifact and its place in history. To foster advanced and thorough research in the field of *DTs* for *TCH*, it is advisable to prioritize the utilization of technical data and *metadata* [158].

Answer to Research Question 2.2 This section explores the methods used for digitalizing *TCH*, with a focus on techniques for capturing data from physical artifacts and converting them into virtual representations. The study investigates the effectiveness of three primary **data acquisition techniques**: *photogrammetry*, *3D laser scanning*, and *sensor-based* approaches [41]. Our focus extended to the key patterns associated with data acquisition techniques and the specific *types of data* collected within the realm of digitalized *TCH*. This perspective aimed to recognize the details of the methodologies employed in capturing and representing cultural artifacts digitally. From this approach, a compelling pattern emerged, highlighting a significant co-occurrence between the utilization of 3D laser scanners and the capture of visual **data type**. This convergence suggests a strategic synergy between the precision offered by 3D laser scanners and the richness of visual data, fostering a comprehensive and detailed representation of cultural artifacts. The prevalence of this co-occurrence underscores the significance of visual fidelity in the digitalization process, emphasizing the importance of capturing subtle details and textures to create a faithful *DT*, as illustrated in Table 3.10. A parallel pattern that emerged in our research involved the frequent co-occurrence of photogrammetry as a data acquisition technique, coupled with the capture of

visual data. Photogrammetry, relying on the principles of extracting information from photographs, emerged as a viable and complementary approach to 3D laser scanning. Combining photogrammetry with visual data further enriches the digitalization process, providing an alternative method that aligns with the imperative of visual accuracy. Within the identified patterns, our research revealed a dominance of visual data as the primary type collected through photogrammetry and 3D laser scanning techniques. However, there were fewer instances of simultaneous utilization of sensor-based data acquisition techniques alongside metadata and technical data types. This underscores the evolving sophistication of *DT* studies at a nascent phase. To conclude, our suggestion supports the adoption of sensors for data acquisition and the integration of metadata into the digitalization process.

3.5.3 Digital Twin Intelligence (RQ2.3)

This subsection explores digital twin operations RQ2.3, “*How does digital twin intelligence leverage advanced operations for tangible cultural heritage?*”. This sub-research question aims to emphasize the operation of digital twins on TCH, supporting the main research question’s objectives regarding the current utilization of digital twin technology. We employed the following parameters to address this question. These include the **activity level of *DT*** assesses the efficiency, reliability, and resilience of current *Digital Twin (DT)* operations in *Tangible Cultural Heritage (TCH)*. The ***DT* functional layer** evaluates the representation, analysis, interaction, and optimization of physical assets within the current *DT* of *TCH*. The selection of these two parameters is driven by their complementary roles in understanding both the operational effectiveness and the functional capabilities of digital twins in the context of cultural heritage. The *activity level* provides insight into how well the *DT* performs in real-world scenarios, while the *functional layer* focuses on how physical assets are digitally represented and managed, ensuring that both technical and practical aspects are fully addressed.

Within the *TCH* sector, *DT* operations are categorized into two main types based on their **activity level**. *Active DTs* not only replicate the physical system but also engage with it in real-time, facilitating control, simulation, and optimization. On the other hand, *passive DT* are representations of physical assets, systems, or processes that primarily concentrate on monitoring [58]. This classification hinges on the degree of activity inherent in the *DT*, serving as a surrogate measure for its maturity level and richness in functionality [226]. Our comprehensive analysis highlights this distinction, showing that a significant proportion of the studies included, specifically 55 ($\approx 51\%$) studies, are categorized as *passive DTs*. The remaining 53 ($\approx 49\%$) studies, constituting the minority, align with the *active DT* type. The distinction between *active* and *passive DTs* serves as a crucial lens through which to assess the maturity and functionality of these digital replicas within the context of *TCH*. The classification is rooted in the activity concept, encapsulating the level of engagement and dynamic interaction facilitated by the *DT* [195].

In the field of *DT* applications, our focus goes beyond simple implementation to explore the **functionality layers**. These layers encompass three essential components: the *presentation layer*, which functions as the interface that enables user interaction with the digital twin and its data, including visualization, user interfaces, and reporting capabilities. The *prediction layer* utilizes machine learning methodologies to anticipate future states, behaviors, or outcomes of the physical

Table 3.11: Co-occurrence of studies on activity level (*RQ2.3.P1*) and functional layer of digital twin (*RQ2.3.P2*).

	Presentation layer	Prediction layer	Prescription layer
Active DT	[49, 89, 91, 96, 114, 181, 198, 247, 262, 284, 285]	[9, 16, 28, 39, 40, 43, 73, 74, 93, 95, 100, 116, 124, 130, 132, 133, 138, 140, 150, 165, 179, 180, 188, 193, 196, 197, 199, 200, 230, 237, 244, 254–256, 263–265, 279, 281, 287, 291, 296]	[1, 7, 41, 72, 98, 166, 220, 278, 288]
Passive DT	[11, 22, 24, 31, 35, 47, 52, 55, 59, 76, 80, 81, 90, 108, 110, 118, 122, 123, 126, 128, 135, 147, 153, 158, 159, 161, 167, 171, 178, 182, 184, 185, 210, 215, 216, 222, 224, 226, 229, 248, 260, 269, 270, 276, 282, 289]	N/A	N/A

system, and the *prescription layer* of a *DT* suggests actions or interventions aimed at enhancing the performance, efficiency, or resilience of the physical system [41, 58]. In our comprehensive study, notably, a substantial emphasis emerges on the *presentation layer*, encapsulating a remarkable 57 ($\approx 53\%$) of the studies analyzed. This layer, positioned at the forefront, primarily concerns itself with the visual representation and portrayal of *DT* data and insights. Following in the hierarchy is the prediction layer, encompassing 42 ($\approx 39\%$) studies, and trailing behind is the prescription layer, with a mere 9 ($\approx 8\%$) studies contributing to this segment. In synthesizing these findings, it becomes apparent that the functional layers of the *DT*, as revealed by our study, are skewed towards the foundational *presentation layer*. While this layer is indispensable for effective communication and understanding, a more comprehensive exploration of the *prediction* and *prescription layers* is warranted to unlock the full potential of *DTs*. The dominance of the *presentation layer* suggests that the current maturity level of *DTs*, especially within the *TCH* domain, may be at an early stage akin to infancy in its development [41].

Answer to Research Question 2.3 In the pursuit of comprehensively understanding the landscape of *DT* application in the *TCH* domain, our exploration has centered around key facets pivotal for assessing the current intelligence status of *DTs*.

After thoroughly analyzing the research, our objective was to identify the pivotal parameters shaping the present scenario in both theory and practice concerning the operation of *DTs* in *TCH*. The findings highlight a notable prevalence of the passive *DT* and the presentation functional layer, as illustrated in Table 3.11. This observation signifies a discernible pattern that characterizes the contemporary landscape of *DT* applications in *TCH*.

Table 3.11 visually encapsulates current trends in *DT* implementation. In particular, the dominance of the passive *DT* and the emphasis on the presentation functional layer align with a focus on

descriptive modeling over the comprehensive exploration of the predictive and prescriptive layers. While revealing the current state, these patterns also spark further inquiries. Slight bonding in the co-occurrence of parameters related to prediction and prescription layers suggests potential growth in *DT* research within *TCH*. The infancy stage indicates ample room for exploration, particularly in embracing dynamic functionalities and advanced modeling techniques. In essence, our study's findings underscore a pivotal moment in the evolution of *DT* applications in *TCH*, signaling the potential for maturation and diversification.

3.6 The Maturity Level of Digital Twin in Tangible Cultural Heritage (RQ3)

This section aims to address two sub-research questions, focusing on the representative model and the implementation level to assess maturity, “*What is the current maturity level of digital twins in the context of tangible cultural heritage?*”. The objective of this research question is to determine the current maturity level of digital twin technology in the tangible cultural heritage domain. To address these objectives, we utilized the following parameters in the sub-research questions **model update frequency**, **twin models**, **model security**, and **verification and validation of DT models**, **implementation stage**, **connectivity mode**, and **computational intelligence** to evaluate the effectiveness of operational efficiency, real-time monitoring and control, and predictive maintenance within *Tangible Cultural Heritage (TCH)* domains.

3.6.1 Models to Assess the Maturity Level of Digital Twin (RQ3.1)

This section is designed to address RQ3.1 *What is a representative model to assess the maturity of digital twin implementations in tangible cultural heritage?* This sub-research question aims to support the main research question's objective by investigating representative models to determine the current maturity level of digital twin technology. In response to this objective, we've employed the following parameters: **model update frequency**, **twin models**, **model security**, and **verification and validation of DT models** because these parameters play essential roles in *DT* development by ensuring that the twin accurately reflects the physical system, operates securely and possesses the computational capabilities.

In the evaluation of *DT* maturity within the *TCH* domain, one critical parameter is **model update frequency** [44]. The intervals at which these updates occur, ranging from *real-time* to *hourly*, *daily*, and *weekly*, significantly influence the maturity level of *DTs*. The decreasing time intervals from *weekly* to *real-time* updates signify the advancement of the *DT* model. Our study's findings reveal a landscape where 35 ($\approx 32\%$) out of the total included studies for weekly updates, 25 ($\approx 23\%$) studies lean towards daily updates, 16 ($\approx 15\%$) studies prefer hourly updates, and 17 ($\approx 16\%$) studies embrace real-time updates. Regrettably, 15 ($\approx 14\%$) studies remain silent on their frequency of updates, lacking clarity on the communication of updates between the physical and digital realms. This discovery emphasizes a compelling need for substantial efforts to push the application of *DTs* in *TCH*, considering the prevalence of studies primarily relying on weekly updates, indicative of a *pre-DT* implementation stage. The absence of a clear update frequency in 15 studies further underscores the infancy of *DT* applications in this domain. The transition from weekly to

real-time updates is not merely a technological advancement but a paradigm shift that allows for a more dynamic and responsive interaction between physical and digital representations. Encouraging a proactive approach to updating practices will boost *DT* applications to higher maturity levels, enhancing their effectiveness in the preservation and exploration of *TCH*.

In the landscape of *DT* implementation, two significant types of **twin models** are evident: *design-time* and *runtime DTs* [120]. Both are integral to the evaluation of *DT* maturity, especially in *TCH* research and applications. *Runtime DTs* offer real-time insights, playing a crucial role in the continuous management and conservation of cultural artifacts. On the other hand, *design-time DTs*, more commonly employed, contribute to the planning and design phases of projects. Our research shows that *design-time DTs* and *runtime DTs* are equally represented in the included studies: each category appears in 54 studies ($\approx 50\%$). This parity indicates that the research community is split between projects focused on planning/design (documentation, 3D reconstruction and analysis) and projects that target operational, sensor-driven behaviors. While design-time *DTs* remain critical for documentation, planning and virtual reconstructions, the equal representation of runtime *DTs* signals a clear emergence of sensor-enabled, real-time monitoring and analytic deployments in *TCH*. Nevertheless, many runtime implementations still rely on limited update frequencies or uni-directional synchronization; advancing runtime capabilities requires improvements in interoperability, scalable data pipelines and automated feedback loops.

In the development and application of *DT* models, safeguarding data emerges as a paramount concern when evaluating the maturity level of *DTs*. These combined measures are crucial for upholding the **security**, privacy, and ethical handling of data within the expansive framework of *DTs* [142]. Within the scope of our study, a noteworthy finding comes to light, revealing that among the 108 included studies, 57 ($\approx 53\%$) were categorized as *non-protected*, while 51 ($\approx 47\%$) fell under the classification of *protected*. This revelation underscores a significant gap in the implementation of security measures within the realm of *DTs*, particularly in *TCH* research. The prevalence of *non-protected* instances implies that security considerations are still in their nascent stages, and the absence of robust protective measures poses challenges to realizing the full potential of *DTs*. Without adequate security protocols in place to shield the communication channels between virtual and physical systems, the effectiveness and reliability of *DTs* are compromised. Addressing this security gap requires a concerted effort to integrate advanced security measures into the *DT* framework [4]. This includes implementing encryption techniques to secure data transmission, establishing stringent access controls to regulate system entry, adopting anonymization practices to protect user identities, and adhering to ethical considerations in data handling. Moreover, cultivating a culture of awareness regarding data security and privacy is crucial for all stakeholders involved in *DT* applications within *TCH* research. In the domain of *DT* applications, the cornerstone of effectiveness lies in the verification and validation of digital twin models, a crucial consideration, particularly within the context of *TCH* research. To navigate and comprehend the multifaceted landscape of *DTs*, we have undertaken a classification of **verification and validation**, dividing it into three distinct dimensions: *model checking*, which entails validating the accuracy, completeness, and consistency of the *DT* model against real-world data and specifications [275]. *Simulation* is utilized to replicate the behavior, performance, and responses of the physical system under various conditions or scenarios. *Testing* involves subjecting the *DT* to real-world conditions or scenarios to assess its performance, reliability, and validity [259]. Our study unravels a subtle distribution within

these classifications, shedding light on the prevailing trends in verification and validation techniques. *Model checking* emerges as a focal point in our analysis, capturing the attention of 46 ($\approx 43\%$) out of the total studies included in our investigation. The prominence of *model checking* signifies a widespread acknowledgment of the importance of rigorously validating the fidelity and accuracy of *DT* representations within the *TCH* research domain. *Simulations*, featuring in 28 ($\approx 26\%$) out of the total studies. *Simulations*, in this context, involve the emulation of real-world scenarios to observe and analyze the *DT*'s responses and behaviors. This dimension adds a layer of dynamism to the understanding of *DT* functionalities, providing insights into how these virtual replicas respond to diverse conditions. In tandem, *testing*, as another facet of model validation, is observed in 14 ($\approx 13\%$) studies within our analysis. In the context of *DTs*, *testing* involves subjecting the virtual replicas to various scenarios to assess their performance, robustness, and reliability. This dimension is crucial for ensuring that *DTs* not only mirror reality but also withstand the complexities of different environmental conditions. However, 20 ($\approx 18\%$) studies, out of the total studies, did not explicitly articulate the model validation techniques. This observation implies a lack of clarity regarding the chosen model validation techniques and suggests that the field is still evolving, with researchers potentially grappling with the intricacies of identifying and articulating the precise purposes of verification and validation in the context of *DTs*.

Answer to Research Question 3.1 In *DT* implementations, evaluating the maturity level is a critical aspect, particularly concerning *TCH* perspectives. Our study focuses on assessing the maturity of *DT* representative models, considering key parameters such as *model update frequency*, *twining model*, *model security*, and *verification and validation of digital twin models*. In the examination of maturity level, we delved into identifying salient parameters consistently emerging in our study, directly linked to the maturation of the *DT* model. As depicted in Table 3.12, a robust correlation was observed between the frequency of weekly model updates and the *DT* during the design phase. This correlation extended to daily updates, exhibiting a frequent association with runtime *DTs*. It is noteworthy that the transition from a weekly update frequency to real-time updates signifies the strength of the model. However, our overall study emphasizes a strong co-occurrence between a weekly update frequency and, a design-time *DT*. This observation directly indicates that the current state of the *DT* representative model is at an early, or infant, stage of maturity.

In the extended investigation of our study, a robust association was found, between weekly updates and non-protected data models. The protected data model shows the correlation with real-time, hourly, and daily updates as encapsulated in Table 3.13. This observation underscores the progression of the *DT* maturity level.

Table 3.14 illustrates the relationship between *model update frequencies* and *the verification and validation (V&V)* methods employed in (*DT*) implementations. As *DTs* progress from foundational to advanced stages, there is a notable shift in update frequency and V&V techniques. In the foundational phase, *DTs* often rely on weekly updates and primarily utilize simulation and model-checking methods for validation. This approach provides a basic level of synchronization between the physical and digital counterparts. As systems mature, the update frequency increases to daily or hourly intervals, incorporating testing methods to enhance accuracy and reliability. *DTs* achieve real-time updates in the most advanced stages, employing comprehensive testing methodologies to ensure immediate and precise reflection of the physical system's state [275]. This progression

Table 3.12: Co-occurrence of studies on model update frequency (*RQ3.1.P1*) and twin models (*RQ3.1.P2*).

	Weekly update	Daily update	Hourly update	Real-time update
Design time DT	[11, 22, 31, 47, 52, 76, 80, 89, 108, 126, 128, 135, 147, 153, 159, 161, 165, 167, 171, 178, 185, 198, 210, 215, 217, 226, 229, 247, 269, 270, 276, 284, 289]	[114, 216]	N/A	N/A
Runtime DT	[73, 96]	[16, 35, 39, 41, 43, 49, 59, 74, 91, 93, 110, 116, 124, 150, 180, 193, 262–264, 279, 281, 291, 296]	[1, 9, 72, 95, 98, 100, 130, 132, 133, 138, 166, 179, 188, 197, 199, 200, 265, 278, 288]	[7, 28, 40, 140, 220, 237, 254–285, 287]

underscores the evolution of the maturity of *DT*, highlighting the importance of frequent data updates and robust V&V processes to maintain the integrity and performance of the system.

3.6.2 Implementation of Digital Twin (RQ3.2)

This section aims to address RQ3.2 : *How can the advancement of digital twin implementation be evaluated to assess its maturity in cultural heritage?* . This sub-research question aims to assess the implementation stage of digital twin technology, supporting the main research question in evaluating its current maturity level in the tangible cultural heritage (TCH) domain. To address this objective, we’ve employed the following parameters: **implementation stage**, **connectivity mode**, and **computational intelligence** because these parameters are crucial in DT development, as they ensure effective communication between the physical and virtual twins, accurately represent the physical system, and integrate the necessary computational capabilities for analysis and simulation. Within the literature, there have been multiple endeavors to establish a framework for gauging the maturity or sophistication levels of *Digital Twins (DTs)* [176]. The initial definition takes a sector-wide perspective, primarily focused on the industry. This categorization delineates the following levels: level 1 or *pre-DT* (virtual system model with an emphasis on technology/technical-risk mitigation); level 2 or *DT* (virtual system model if the physical twin exists); level 3 or *adaptive DT* (virtual system model of the physical twin with adaptive UI); level 4 *intelligent DT* (virtual system model of the physical twin with adaptive UI and reinforcement learning) [132, 249]. From this definitional framework, our exploration delves into the **implementation stage**. This investigation aims to measure the maturity level of *DT* applications, specifically in the context of *TCH*. Our study presents an in-depth picture, revealing that within the spectrum of digital twins, *DT*, and adaptive *DT*, *DT* holds the most influence, dominating the landscape with 59 ($\approx 55\%$) out of the 108 included studies. This is followed by the *pre-DT* category, accounting for 40 ($\approx 37\%$) studies, while *adaptive*

Table 3.13: Co-occurrence of studies on model update frequency (*RQ3.1.P1*) and Model security (*RQ3.1.P3*).

	Weekly update	Daily update	Hourly update	Real-time update
Protected	[73]	[16, 35, 39, 41, 49, 74, 93, 110, 116, 124, 150, 193, 262–264, 279, 281, 291, 296]	[1, 9, 95, 98, 132, 133, 138, 166, 179, 188, 197, 199, 200, 265, 278, 288]	[7, 28, 40, 100, 140, 220, 230, 237, 244, 254–256, 285, 287]
Non-protected	[11, 22, 31, 47, 52, 73, 76, 80, 89, 96, 108, 126, 128, 135, 147, 153, 159, 161, 165, 167, 171, 178, 185, 198, 210, 215, 217, 226, 229, 247, 270, 276, 284, 289]	[43, 59, 91, 114, 180, 216]	[72, 130, 196]	

digital twin claim a smaller share, represented in only 9 ($\approx 8\%$) studies. Notably, the intelligent DT category remains unexplored, as no studies have delved into its maturity level, presenting an area ripe for further investigation. This signals an infancy stage, calling for concerted efforts to propel *DT* applications to higher levels of sophistication. The unexplored realm of intelligent *DTs* provides an avenue for future research, presenting an exciting opportunity to advance the capabilities of *DTs* within the rich tapestry of *TCH*.

DTs tailored for *TCH*, the profound impact of **connectivity modes** on the orchestration of data flow between the physical entity and its virtual counterpart is paramount. Let's delve into an in-depth exploration of *bidirectional*, *unidirectional*, and *automatic* connectivity within the unique context of *DTs* designed for *TCH* preservation and exploration [288]. The selection of a specific connectivity mode emerges as a pivotal factor, serving as a critical determinant in the comprehensive assessment of the *DT* maturity level [249]. From this perspective, our comprehensive studies embark on an exploration of the existing connectivity landscape, which includes 108 meticulously selected studies. The findings shed light on the prevalent connectivity status within this domain. Notably, 51 ($\approx 47\%$) out of the total studies predominantly employ *unidirectional* connectivity characterized by the singular flow of data in either direction, from the physical entity to the *DT* or vice versa. This *unidirectional* approach is instrumental in scenarios where a one-way exchange of data suffices for the objectives of monitoring or control. Conversely, *bidirectional* connectivity, denoting a dynamic and reciprocal exchange of data between the physical entity and its digital representation in both directions, is evident in 47 ($\approx 44\%$) studies. This mode facilitates real-time interactions, feedback loops, and a synchronized exchange of information, enriching the *DT* capabilities. Moreover, the examination reveals that *automatic* connectivity, where the *DT* autonomously initiates, establishes, and maintains connections with the physical entity without requiring manual intervention, is featured in 5 ($\approx 5\%$) studies. This autonomy streamlines processes, enhancing op-

Table 3.14: Co-occurrence of studies on model update frequency (RQ3.1.P1) and Verification and validation of DT models (RQ3.1.P4).

	Weekly update	Daily update	Hourly update	Real-time update
Simulation	[11, 47, 73, 108, 126, 135, 147, 159, 198, 210, 215, 229, 270, 276, 284, 289]	[114, 263, 291]	[1, 9, 95, 188, 197]	[100, 256]
Model-checking	[22, 89, 96, 153, 165, 178, 217, 226]	[16, 43, 74, 91, 93, 110, 116, 124, 150, 180, 193, 262, 264, 279, 296]	[72, 130, 133, 166, 179, 199, 200, 265, 278, 288]	[31, 40, 140, 220, 244, 285, 287]
Testing	[31]	[35, 41, 49, 281]	[98, 132, 138]	[7, 28, 230, 237, 254, 255]

erational efficiency and reducing the need for constant human oversight. However, it's noteworthy that in 5($\approx 5\%$) studies, connectivity between the physical and virtual entities is not explicitly demonstrated.

In digital twin advancement, **computational intelligence** plays a significant role in evaluating the maturity of *DTs*. Computational capabilities combine *statistical modeling* to comprehend the behavior of the system using historical data, *predictive modeling* to anticipate future behavior and outcomes, and *adaptive behavior* to adapt and learn autonomously from the environment [213]. The results of our study highlight the significant emphasis placed on predictive modeling, which was utilized in 47 ($\approx 44\%$) out of 108 surveyed studies. Statistical modeling is used in 36 ($\approx 33\%$) studies, while adaptive behavior is found in only 10 ($\approx 9\%$) studies. In particular, 15 ($\approx 14\%$) studies do not explicitly outline the modeling techniques used within computational intelligence. These discoveries have significant implications, indicating a dominant tendency in *DT* implementation that places a strong emphasis on statistical modeling. This observation initiates a thorough evaluation of the current maturity level of *DTs* at the initial stage, particularly within the *TCH* domain.

Answer to Research Question 3.2 Assessing the maturity level of *DTs* in the *TCH* domain involves a multifaceted examination, encompassing *implementation stages*, *computational intelligence* and *connectivity modes*. This task is to identify the current implementation stage, evaluate the connectivity mode and computational capacity, which are integral elements in assessing the sophistication and effectiveness of *DT* applications. In measuring the implementation stage, well-established categories are employed, delineating the journey from *pre-DT* to *intelligent DT*. Each stage aligns with specific connectivity modes such as unidirectional, bidirectional, and automatic, showcasing the evolution in information flow as the implementation stage progresses. The correlation between *implementation stages*, *connectivity modes*, and *computational intelligence* illustrates that as the implementation stage advances from *pre-DT* to *intelligent DT*, the strength of connectivity also advances from unidirectional to automatic information flow, while computational intelligence

Table 3.15: Co-occurrence of studies on stage of digital twin implementation (*RQ3.2.P1*) and connectivity mode (*RQ3.2.P2*).

	Pre-digital twin	Digital twin	Adaptive digital twin	Intelligent digital twin
Unidirectional	[11, 22, 47, 59, 76, 80, 81, 89, 90, 108, 122, 123, 126, 128, 135, 147, 158, 159, 161, 167, 171, 178, 182, 185, 210, 215–217, 229, 247, 248, 260, 269, 276, 284, 289]	[31, 43, 52, 55, 73, 96, 153, 165, 181, 184, 198, 226, 270, 279]	N/A	N/A
Bidirectional	N/A	[16, 28, 35, 39, 40, 49, 72, 74, 91, 93, 95, 98, 100, 110, 114, 116, 124, 130, 133, 138, 140, 150, 179, 180, 188, 193, 196, 197, 200, 220, 230, 237, 244, 254–256, 262–265, 281, 285, 287, 296]	[1, 98, 132, 199, 220]	N/A
Automatic	N/A	[166]	[7, 41, 278, 288]	N/A

advances from statistical modeling to adaptive behavior. The prevailing connectivity paradigm in *TCH* leans heavily towards unidirectional modes, predominantly featuring *DT* stages. This underscores the foundational maturity level, prompting a call for exploration into advanced connectivity modes like bidirectional and automatic communication, as well as predictive modeling and adaptive behavior. We delved into prominent patterns, evident in the co-occurrence displayed in the provided Table 3.15 and 3.16 consecutively.

3.7 Discussion

this thesis provides an in-depth exploration of the rapidly evolving applications of *Digital Twin (DT)* technology in *Tangible Cultural Heritage (TCH)*, emphasizing its transformative potential from a user perspective, the enabling technologies driving its adoption, and an evaluation of its current maturity level [58, 87].

Figure 3.2 presents a tri-dimensional, taxonomy-informed conceptual framework for *DT* applications in *TCH*, addressing RQ1–RQ3 along three dimensions: user-centric applications, enabling technologies, and maturity and operations. Relationship strengths (strong, moderate, weak) ,determined based on the frequency and consistency of evidence reported in the reviewed literature.. User-centric applications cover artwork types, purposes (preservation, conservation, restoration), and *DT* techniques, showing preservation relies on photogrammetry and laser scanning, conserva-

Table 3.16: Co-occurrence of studies on stage of digital twin implementation (*RQ3.2.P1*) and Computational intelligence (*RQ3.2.P3*).

	Pre-digital twin	Digital twin	Adaptive digital twin	Intelligent digital twin
Statistical modeling	[13, 22, 47, 59, 76, 90, 126, 147, 158, 159, 161, 167, 178, 182, 210, 215–217, 224, 229, 247, 276, 284, 289]	[31, 35, 49, 52, 55, 91, 116, 153, 184, 198, 226, 270]	N/A	N/A
Predictive modeling	[89]	[16, 28, 39, 40, 43, 72–74, 93, 95, 96, 100, 110, 124, 130, 133, 138, 140, 150, 165, 166, 180, 188, 193, 196, 200, 230, 237, 244, 254, 256, 263–265, 265, 279, 281, 285, 287, 296]	[132, 199]	N/A
Adaptive behaviors	N/A	[197, 255, 262]	[1, 7, 41, 98, 220, 278, 288]	N/A

tion uses sensor-based monitoring, and public-oriented applications favor image-based visualization while immersive VR/AR is limited. Enabling technologies include data acquisition, 3D modeling, and visualization platforms, with design-time models relying on photogrammetry/laser scanning, runtime monitoring on sensors, and platforms like Unity Reflect enabling immersive experiences. Maturity and operations capture lifecycle stages, twin models, functional layers, connectivity, update frequency, security, V&V, and computational intelligence, revealing early-stage dominance with presentation layers, unidirectional connectivity, limited predictive capabilities, and intelligence evolving from statistical models to adaptive behaviors in advanced *DTs*. **User-Centric Applications of DT:** *DT* technology in *TCH* has fundamentally transformed how users interact with and experience cultural assets. From the user’s perspective, *DTs* serve as powerful tools for public engagement, educational outreach, and research facilitation. Immersive technologies such as Virtual Reality (VR) and Augmented Reality (AR) enrich user interactions, offering immersive explorations of heritage artifacts and sites [158, 167]. However, recent studies reveal that static 3D models remain predominant, while fully interactive *DT*-driven VR/AR experiences are rare, highlighting a significant gap between research and practice [237, 276]. This gap highlights the opportunity to expand immersive platforms to make *TCH* experiences more dynamic, inclusive, and engaging, particularly for younger audiences and students. In practice, the Louvre Abu Dhabi’s XR project




Figure 3.2: Taxonomy-informed conceptual framework for *DT* applications in tangible cultural heritage, structured along three complementary dimensions: user-centric applications, enabling technologies, and maturity levels, along with their interdependencies. Arrow thickness indicates relationship strength: thick lines = strong, medium lines = moderate, thin lines = weak.

and the Virtual Rome initiative demonstrate how *DT*-driven *VR/AR* can scale heritage education and tourism, but these remain isolated cases rather than systemic adoption [63, 77]. Additionally, *DTs* offer researchers a risk-free method to analyze artifacts, while educators and students gain unprecedented access to interactive learning resources [89]. Nevertheless, the limited focus on these user groups in current research underscores the need to broaden educational and accessibility-driven applications of *DTs*, ensuring inclusivity and engagement for diverse audiences.

Enabling Technologies: The development and deployment of *DTs* for *TCH* rely on advanced enabling technologies, including data acquisition methods, modeling tools, and platforms. While 3D laser scanning and photogrammetry dominate for their precision and efficiency, sensor-based methods remain underutilized. Integration of real-time environmental monitoring through sensor data could greatly enhance *DTs*, enabling dynamic updates and proactive conservation strategies [180]. Current *DT* implementations heavily emphasize the presentation layer (visualization), but predictive and prescriptive layers remain underdeveloped. Progress in these areas is vital for unlocking the full potential of *DTs* as dynamic, intelligent systems capable of real-time monitoring, simulation, and intervention. Recent findings confirm that most *TCH*-related *DTs* are still “unidirectional,” synchronizing data from physical to digital, but lacking feedback loops to inform real-world interventions [61]. Advancing connectivity and cyber-physical intelligence could enable *DTs* to transition from passive representations to adaptive, interactive systems that respond dynamically to environmental and operational changes [213]. Moreover, the limited application of runtime *DTs* within the domain reflects the nascent maturity of the field. While design-time *DTs* play an essential role in planning and strategizing cultural preservation, runtime *DTs* hold immense potential for real-time intervention and dynamic operational control. For instance, the Notre-Dame *DT* project integrates

sensor-driven runtime monitoring for structural risk prediction, demonstrating the practical feasibility of real-time updates for heritage resilience.

Evaluating Digital Twin Maturity Levels: The maturity of *DT* applications in tangible cultural heritage is still in its infancy. Most implementations align with foundational *DT* or pre-*DT* stages, characterized by static digital models and limited real-time interaction. Adaptive and intelligent *DTs*, which incorporate real-time updates, predictive modeling, and autonomous behaviors, remain largely unexplored. Weekly or daily model updates, rather than real-time synchronization, dominate heritage *DT* implementations, confirming the observations of [138, 244, 279] that maturity in this field lags far behind industrial *DT* practice. Significant gaps in data security, simulation capabilities, and real-time interactivity hinder the evolution of *DTs*. Addressing these challenges through investments in advanced technologies, robust frameworks, and interdisciplinary collaboration is crucial for enhancing *DT* maturity. Such advancements would enable *DTs* to transition from basic visualization tools to comprehensive systems capable of supporting predictive and prescriptive functionalities for conservation and management. To support this transition, leveraging advanced platforms such as Unity Reflect, Unreal Engine, or cloud-based solutions like AWS IoT TwinMaker can provide scalable infrastructure for managing complex data and real-time synchronization in heritage *DT* projects [132, 249]. Furthermore, fostering collaborative ecosystems that involve stakeholders from technology, heritage conservation, and academia will accelerate the adoption and sophistication of *DTs*.

A further step is the adoption of structured maturity assessment models tailored to TCH. While industrial *DTs* already employ multi-level frameworks (e.g., NASA’s five-level *DT* maturity model), heritage-focused assessments remain fragmented. Future work should establish evaluation models that consider parameters such as update frequency, bidirectional connectivity, metadata integration, and user inclusivity. Embedding maturity models into heritage *DT* workflows would allow the field to progress systematically toward adaptive, intelligent, and sustainable *DTs* for cultural heritage. This approach ensures not only academic contribution but also direct impact on real-world practices in conservation, education, and tourism.

3.7.1 Historical context and evolution *DT*

DT technology, originally developed for industrial and engineering sectors, has become a key tool to improve the monitoring, conservation, preservation, dissemination, and management of *TCH* assets [58, 226]. Recent research highlights the importance of standardized UX evaluation methodologies, such as ISO 9241-210, in assessing immersive heritage platforms [10]. These standards help ensure that XR and metaverse applications are not only engaging but also accessible and effective for diverse audiences. Furthermore, XR-based virtual tourism frameworks provide sustainable alternatives to physical site visits, reducing wear on fragile heritage assets while expanding global access to cultural experiences. By creating dynamic digital replicas of physical entities such as artifacts, monuments, and heritage sites, *DTs* enable real-time data monitoring, simulation, and interaction, fostering innovative methods to protect these invaluable assets from degradation, natural disasters, and the passage of time [166]. The transition from traditional heritage documentation techniques (drawings, casts, photographs) to digitally integrated workflows underscores the shift toward data-driven preservation, powered by IoT, AI, machine learning, 3D scanning, and extended reality (XR) [54]. Recent studies confirm that such integrations are increasingly common in real-




Figure 3.3: Timeline of *DT* evolution in cultural heritage (2002–2025), highlighting milestones from conceptual origin to intelligent and federated *DT* frameworks. Adapted from [7, 10, 51, 87, 97, 101, 122, 123, 169, 239].

world projects such as the Notre-Dame restoration, *DT*-based monitoring of church artworks using IoT, and XR-supported museum education platforms [61, 97].

Figure 3.3 presents a chronological overview of the evolution of *DT* technology from its conceptual origin in 2002 to advanced applications in cultural heritage by 2025. The timeline highlights major milestones, including the initial development of *DT* concepts for industrial systems (2002–2010), early adoption in cultural heritage contexts (2016), integration with virtual museums and XR platforms (2019), and the incorporation of AI and IoT for heritage monitoring (2022). Recent developments emphasize immersive XR and metaverse-based experiences (2023), intelligent *DTs* and advanced modeling (2024), and federated *DT* frameworks with sustainability considerations (2025). This evolution reflects the transition from static digital representations to adaptive, intelligent, and collaborative *DT* systems for cultural heritage.

3.8 Open Challenges and Future Directions

This section highlights the open research challenges identified through the analysis of *Digital Twin (DT)* technology in *Tangible Cultural Heritage (TCH)* and outlines future research directions to address these gaps. These challenges span technological, methodological, legal, and practical aspects.

3.8.1 Open Challenges from the Analysis

Enhanced Data Interoperability: A major challenge is the lack of interoperability across heterogeneous data sources and systems [262]. Cultural heritage sites produce diverse datasets, including 3D scans, IoT sensor data, environmental readings, historical records, and architectural blueprints. Current *DT* implementations rarely integrate these datasets within a unified framework, limiting cross-institutional collaboration. The absence of common standards not only prevents data sharing but also increases costs when institutions rely on incompatible proprietary solutions. Developing open standards, metadata schemas, and interoperable protocols is therefore essential to facilitate seamless data exchange and enable multi-stakeholder collaboration [41, 255].

Comprehensive Data Integration: While most *DTs* capture structural and geometric data,

contextual information (e.g., cultural significance, oral histories, intangible aspects) is often neglected [38, 256]. This limitation reduces the authenticity of digital replicas and undermines their educational and interpretive value. A richer integration of metadata, technical specifications, and historical narratives would enable more holistic and meaningful representations. For example, integrating environmental monitoring data with historical context could guide restoration strategies by balancing material preservation with cultural symbolism [72, 166].

Advanced Modeling Techniques: Existing *DT* models often focus on visualization, while predictive and prescriptive capabilities remain underdeveloped [288]. Hybrid approaches that combine data-driven methods (e.g., machine learning) with physics-based models (e.g., structural behavior, material decay) can better capture the complexities of historical structures and artifacts. Recent works, such as AI-driven monitoring of church artworks and *DT*-assisted Notre Dame restoration projects, demonstrate the value of predictive modeling for conservation. Such hybrid models can support decision-making by simulating degradation processes, forecasting risks, and prioritizing interventions [72].

Immersive and Scalable Applications: Although Virtual Reality (VR) and Augmented Reality (AR) have been piloted in cultural heritage, most applications remain limited to static visualization [216]. Expanding immersive, interactive, and scalable *DTs* could significantly enhance public engagement, tourism, and education. For example, VR environments can replicate inaccessible or damaged heritage sites for remote exploration, while AR can enrich onsite visits with contextual overlays [80, 180]. However, accessibility is uneven: high-end immersive applications require advanced devices, leaving mobile-based, web-friendly solutions underdeveloped. Scaling *DTs* across devices and platforms is therefore critical to ensure inclusivity.

Robust Security and Ethical Concerns: As *DTs* increasingly store sensitive cultural heritage data, cybersecurity and ethical issues become paramount [118, 262]. Unauthorized access, malicious alterations, or digital piracy could compromise authenticity. Blockchain can ensure data integrity by maintaining immutable change logs, while encryption safeguards data transmission and storage [288]. In addition, governance and ethical frameworks are required to address questions of digital ownership, copyright, and data sovereignty, particularly when heritage data crosses national and institutional boundaries [285].

3.8.2 Future Directions

Building on these challenges, future research must adopt more practical, interdisciplinary, and forward-looking approaches to *DTs* in cultural heritage.

First, **integration of AI, IoT, and cloud computing** should move beyond visualization to enable predictive analytics, adaptive conservation, and real-time decision support. Case studies such as Notre Dame’s *DT*-enabled restoration and IoT-based church monitoring demonstrate that these approaches are already feasible but require broader validation and standardization.

Second, the adoption of **Federated Digital Twins (FDTs)** offers a promising direction. Unlike centralized *DTs*, *FDTs* enable secure, decentralized collaboration across multiple institutions. They address interoperability and privacy challenges by allowing distributed data processing while ensuring local data sovereignty. *FDT* frameworks could standardize formats, governance protocols, and data sharing agreements between heritage institutions, making *DT* solutions more scalable and sustainable.

Third, **maturity assessment frameworks** are urgently needed. Current *DT* implementations lack structured evaluation, hindering progress from static models toward intelligent, adaptive systems. Future research should establish benchmarks that assess update frequency, interactivity, bidirectional connectivity, and integration of contextual metadata.

Finally, **inclusive and accessible *DT* design** should be prioritized. This involves creating lightweight, mobile-first, and web-based solutions for diverse audiences, including tourists, educators, and scholars. By lowering technological barriers, *DTs* can support equitable access to cultural heritage across geographic, social, and economic boundaries. In addition, XR-based virtual tourism frameworks, as highlighted by Anwar et al.[10], offer scalable and inclusive access to cultural heritage. These immersive environments allow users to explore heritage sites remotely, reducing physical degradation while enhancing educational and emotional engagement. Future *DT* designs should incorporate such frameworks to support equitable access across geographic and socioeconomic boundaries. In general, advancing toward interoperable, intelligent, and federated *DT* ecosystems can ensure that the preservation of cultural heritage is not only technologically advanced, but also ethically grounded, accessible, and sustainable. This trajectory positions *DTs* as vital tools for safeguarding heritage while enriching education, public engagement, and research in the digital era.

3.9 Conclusion

This study presents a comprehensive review of Digital Twin (*DT*) applications in Tangible Cultural Heritage (*TCH*), structured around user-centric applications, enabling technologies, and maturity levels. Findings reveal that most *DTs* remain in early-stage development, relying on static models and unidirectional data flows. Interactive features, real-time updates, and adaptive intelligence are underutilized. The lack of standardized maturity frameworks and limited integration of metadata and sensor-based data hinder progress. To address these gaps, we propose adopting Federated Digital Twins (*FDTs*), cloud-based platforms, and inclusive design strategies that support diverse users and devices. Embedding structured evaluation models and nurturing interdisciplinary collaboration will allow *DTs* to evolve into intelligent, scalable, and ethically grounded systems. These advances will improve the preservation, education and public participation of cultural heritage in the digital era. In particular, our tri-dimensional framework provides a roadmap that can guide museums, heritage authorities, and real-world practices in designing scalable and interoperable *DT* solutions while ensuring their evolution toward inclusive, sustainable, and high-impact applications. Beyond the academic contribution, this study offers an instrument to position cultural heritage's current practices with respect to the state-of-the-art, and actionable recommendations for researchers, policymakers, and practitioners. By bridging synthesis with practice, this study highlights not only where research has advanced but also where implementation must progress to achieve long-term sustainability in cultural heritage.

The findings of this systematic analysis reveal that most Digital Twin implementations in cultural heritage remain fragmented and lack intelligent, collaborative capabilities. These limitations motivate the development of the AI-driven Federated Digital Twin framework presented in the following chapter.

Chapter 4

AI-Driven Federated Digital Twin for Cultural Heritage

Overview of Chapter

This chapter presents the second major implementation-oriented contribution of the dissertation, building directly upon the foundational concepts of Digital Twins (Chapter 2), the systematic analysis of their applications in cultural heritage (Chapter 3). This chapter introduces a simulation-based prototype of an AI-driven Federated Digital Twin (FDT) architecture designed to support decentralized, privacy-preserving, and intelligent emergency response across distributed cultural heritage institutions. The chapter details a multi-layered FDT framework that integrates local digital twins, federated learning, and agent-based modeling to enable real-time hazard detection, adaptive evacuation coordination, and ethical data governance. Simulation experiments conducted using NetLogo demonstrate the framework’s effectiveness in managing fire emergencies across two museum environments, achieving a 93.8% evacuation success rate and 88.8% routing efficiency across simulation runs. Key contributions include a federated architecture that preserves institutional autonomy and data privacy, semantic interoperability using CIDOC CRM for cross-institutional coordination, simulation-based validation of emergency scenarios with heterogeneous agent behaviors, and mathematical modeling of federated learning, decision logic, and optimization objectives. This chapter incorporates, with editorial adjustments, the core content of the article “AI-Driven Federated Digital Twins for Scalable Emergency Response in Cultural Heritage Sites”, accepted for publication in the ACM Journal on Computing and Cultural Heritage (JOCCH). Figures, tables, and simulation results have been reformatted to align with the dissertation’s structure, while maintaining the integrity of the original scientific contributions. References have been harmonized with the thesis-wide bibliography, and cross-references have been added to ensure coherence with other chapters.

4.1 Introduction

The preservation of cultural heritage in the digital age requires intelligent, adaptive, and secure technologies that support both the physical safeguarding of artifacts and the operational resilience of cultural institutions [15, 134]. From historic monuments and museum collections to intangible traditions and visitor experiences, cultural heritage is increasingly exposed to threats, such as natural disasters, human-induced hazards, and climate-related events. These risks are further am-

plified in multisite institutions or heritage clusters, where coordination, real-time situational awareness, and data-driven decision-making are essential for ensuring safety, continuity, and accessibility [12, 79, 258].

Recent years have witnessed the evolution of Digital Twin (DT) technology as a transformative approach for real-time monitoring, system simulation, and adaptive control [46, 78]. DTs are virtual representations of physical systems that continuously receive data from sensors and provide actionable insights through analysis and simulation. To extend this concept to large-scale, distributed, or multi-institutional settings, the notion of Federated Digital Twins (FDTs) has emerged. FDTs enable multiple DTs, each representing a specific physical entity, subsystem, or institution, to operate autonomously while collaborating through a decentralized framework. This allows for horizontal integration (across different institutions or sites) and vertical integration (across management levels), while maintaining local autonomy and data privacy [2, 208].

Despite the growing maturity of DT and FDT architectures in domains such as manufacturing, logistics, and critical infrastructure, their adoption in the cultural heritage sector remains limited and largely underexplored. Heritage institutions face domain-specific challenges that differentiate them from industrial or urban environments. These include fragmented governance structures, the immovability and fragility of assets, strict data protection and ethical considerations, and highly variable technical infrastructure across sites [2, 268]. Moreover, emergency scenarios in cultural heritage settings are uniquely complex, as they require protecting both people and irreplaceable artifacts, often within historically significant or structurally constrained environments. These challenges necessitate a new generation of digital twin architectures that are not only privacy-preserving and interoperable but also context-aware, minimally invasive, and ethically aligned with heritage conservation principles.

Existing DT implementations in the cultural heritage domain predominantly focus on static digital documentation, such as 3D scanning of monuments or condition monitoring of artifacts [174]. Although valuable, these applications often lack real-time adaptability, predictive capabilities, and mechanisms for inter-institutional coordination, particularly in emergency response and operational resilience. Most current digital twin initiatives in heritage contexts do not address the complexity of distributed real-time decision-making or hazard mitigation across multiple institutions [258]. This represents a critical research and practice gap that hinders the broader adoption of DT technologies in cultural preservation.

This study addresses this scientific and operational gap by proposing an AI-driven Federated Digital Twin (FDT) architecture specifically designed for scalable, privacy-aware emergency response in interconnected cultural heritage environments. The framework enables decentralized monitoring, hazard prediction, localized decision-making, and federated coordination among distributed heritage institutions, without requiring centralized data sharing. To evaluate the framework, we implement agent-based evacuation simulations replicating fire emergencies across two virtual museums. These simulations assess evacuation success rates, routing efficiency, and cross-site collaboration under dynamic, multi-hazard scenarios. By integrating concepts from AI, federated learning, simulation modeling, and cultural heritage informatics, this research introduces a novel approach to enhancing resilience, situational awareness, and ethical coordination in multi-site cultural heritage contexts. The proposed architecture lays the foundation for future systems that support both tangible and intangible heritage protection through secure, scalable, and intelligent infrastructure. To contex-

tualize this gap, the following section reviews the technical evolution of digital twin systems and highlights the challenges of adapting federated frameworks to the unique conditions of cultural heritage environments.

4.2 Background

The concept of Digital Twin (DT) technology, first introduced by Grieves and Vickers, involves creating a digital representation of a physical entity that allows for the monitoring, simulation, and analysis of real-time data from the physical world [48, 86]. A DT system fosters continuous feedback between the physical asset and its digital model, allowing organizations to forecast potential failures, improve performance, and make data-driven decisions. Using IoT sensors, machine learning algorithms, and cloud computing, DTs capture and process real-time information, providing substantial benefits in areas such as predictive maintenance, improved operational efficiency, and optimized resource management [8, 241]. While the evolution from traditional Digital Twins (DTs) to Federated Digital Twins (FDTs) systems has unlocked new capabilities in distributed intelligence, most documented applications are tailored to highly integrated domains such as smart cities, industrial automation, and healthcare [66, 207]. These environments typically feature standardized infrastructure, centralized control, and unified data governance conditions that do not reflect the operational diversity of cultural heritage institutions. Museums and heritage sites are often decentralized, governed by heterogeneous policies, and limited in terms of technical infrastructure and funding [79, 174]. Moreover, emergency response in heritage contexts involves protecting not only human lives but also immovable and historically significant artifacts, introducing complex ethical, logistical, and preservation challenges. These unique constraints demand a specialized FDT architecture capable of enabling real-time, coordinated decision-making across loosely connected, risk-prone, and privacy-sensitive cultural environments.

4.2.1 Standalone Digital Twin vs Federated Digital Twin

The concept of a Federated Digital Twin (FDT) represents an advanced and interconnected approach to digital twin technology, where multiple digital twins are integrated within a cohesive framework [27, 207]. Unlike traditional Digital Twins (DTs), which typically operate in isolation and focus on a specific asset or domain, FDTs facilitate data sharing, communication, and collaborative decision-making across diverse systems or locations. This integration allows for a more holistic and system-wide approach to monitoring, analysis, and optimization, which is particularly valuable in interconnected environments such as smart cities, industrial networks, or cultural heritage clusters. In this sense, FDTs can be seen as a natural evolution of traditional DTs, expanding their scope through federated architectures that enable efficient data flow and coordination between distinct digital twins [105, 232]. The primary distinction between DTs and FDTs lies in their architecture. Traditional DTs function independently, focusing on a particular entity, be it machinery, vehicles, or infrastructure, within a controlled environment [292]. These digital replicas are used primarily in sectors where monitoring specific asset data (e.g., temperature, pressure, and operational performance) is critical for predictive maintenance and optimization [232]. For example, DTs in manufacturing environments allow for the monitoring of individual machines to predict potential breakdowns and improve productivity. However, the asset-centric nature of DTs limits their

application to isolated use cases. This limitation is particularly evident in scenarios that require coordination between multiple DTs or systems, as traditional DTs lack the interoperability required to seamlessly integrate with other systems [20, 207]. FDTs address this limitation by creating a federated framework that enables interconnected DTs to operate collectively. FDTs leverage an interoperable platform that facilitates data flow and decision-making across different DTs, all while maintaining the operational independence of each twin. This architecture is especially beneficial in domains like urban infrastructure management, where coordinating multiple assets ensures optimal system performance. For instance, in a smart city scenario, an FDT model can integrate DTs representing traffic systems, utility grids, and emergency response services, fostering a comprehensive, city-wide approach to data-driven management [207, 292]. Interoperability is the key to FDT systems, as it enables seamless communication and data exchange between the digital twins, allowing for coordinated responses and optimizations across various domains [129, 155]. The architectural evolution and comparison between traditional DTs and FDTs is illustrated in Figure 4.1 .




Figure 4.1: Comparison of Digital Twin and Federated Digital Twin across key technical and operational aspects.

4.2.2 Related Work

Digital Twins (DTs) have been widely adopted in various industries, including manufacturing, healthcare, and smart cities, to enable real-time monitoring, predictive maintenance, and simulation-based decision-making [86, 241]. Traditional DTs focus primarily on isolated assets, limiting their ability to coordinate across distributed environments. Recent studies have explored DT applications

in the preservation of cultural heritage, using 3D modeling, IoT-based monitoring, and AI-driven analytics to improve site management and visitor engagement [32, 173]. However, these approaches lack the scalability and interoperability to manage multiple interconnected heritage sites [32, 207]. To overcome the limitations of standalone DTs, Federated Digital Twins (FDTs) have emerged as a promising paradigm, enabling decentralized data exchange and collaborative decision-making between multiple entities [208]. FDTs have been successfully deployed in smart cities, logistics, and energy management, where distributed systems must coordinate in real time [67].

In the cultural heritage domain, recent efforts have explored networked DTs for multi-site monitoring, but challenges related to interoperability, secure data sharing, and adaptive learning remain largely unaddressed [72]. Federated Learning (FL) has been integrated into DT architectures to enhance predictive analytics while preserving data privacy [207]. Studies in disaster response and public safety demonstrate that FL can enable real-time situational awareness, optimize evacuation strategies, and support decision-making in dynamic environments [5]. Existing research highlights the potential of FL to optimize emergency management in urban infrastructure [177], but its application in heritage sites remains underexplored. Integrating FL with FDT can provide a decentralized yet coordinated approach to managing emergency scenarios across multiple sites. The use of DTs for risk assessment and disaster preparedness in cultural heritage sites has gained increasing attention. Previous studies have used DT-based models for monitoring structural health, fire risk analysis, and crowd management [6, 67]. Although these efforts demonstrate the value of DTs in preserving heritage assets, they often lack real-time coordination across multiple locations, limiting their effectiveness in large-scale emergency scenarios. A federated approach could enable synchronized evacuation plans and dynamic risk assessment in geographically distributed sites [119]. Despite the growing interest in digital transformation for heritage preservation, the integration of federated learning within digital twin systems remains in its early stages. Prior efforts primarily focus on visualization, environmental monitoring, or isolated modeling, with little attention to dynamic coordination across multi-institutional settings. Building on the challenges outlined in 4.2, our approach introduces a novel AI-enhanced FDT framework that combines federated learning and real-time simulation for emergency preparedness. This architecture is specifically designed to enhance visitor safety, artifact protection, and operational decision-making under distributed, privacy-sensitive conditions.

4.3 Proposed Federated Digital Twin Framework

The proposed Federated Digital Twin (FDT) framework is a multi-layered architecture designed to enhance real-time monitoring, analytics, and emergency response coordination across interconnected cultural heritage sites, as illustrated in Figure 4.2. While this architecture is primarily applied to museum networks, it has significant potential for broader applications in cultural heritage management. Cultural heritage encompasses not only physical structures and artifacts but also intangible elements such as rituals, storytelling, visitor experiences, and collective memory. These diverse elements are increasingly vulnerable to natural disasters, climate impacts, and human-induced threats, which necessitate adaptive, ethical, and scalable digital infrastructure.

The architecture is composed of five integrated layers that work in harmony to create a dynamic, intelligent, and secure system tailored to the needs of decentralized cultural heritage institutions.

At the foundation lies the **Physical Layer**, which represents the real-world environment of each museum or heritage site. This layer includes sensing technologies such as motion detectors, environmental monitors, and IoT devices that continuously gather data on visitor movement, temperature, humidity, air quality, and potential hazards like fire or structural strain. These real-time data streams are critical for both risk assessment and operational awareness. Building directly upon the Physical Layer, the **DT Layer** models each museum using a Local Digital Twin (LDT). Each LDT performs localized data analysis, simulates emergency scenarios, and supports real-time decision-making specific to the institution's layout, artifact sensitivity, visitor density, and historical context. For example, an LDT can recommend optimized evacuation routes, predict potential bottlenecks, and prioritize actions to protect fragile exhibits. This layer supports autonomy and adaptability by allowing each institution to respond to evolving conditions without external dependencies. Interfacing with the DT Layer, the **Aggregation Layer** serves as the system's coordination engine. It collects and harmonizes outputs from individual LDTs into a unified knowledge base, enabling federated learning and continuous optimization. This aggregation enables institutions to benefit from shared intelligence while preserving their data autonomy. Furthermore, it ensures interoperability through standardized data formats, analytics models, and decision-making protocols across diverse heritage contexts. To strengthen semantic consistency and interoperability across heterogeneous cultural heritage institutions, the framework can incorporate established ontologies such as the CIDOC Conceptual Reference Model (CIDOC CRM) [68]. CIDOC CRM provides a formalized ontology for representing cultural heritage entities, events, and relationships, thereby enabling harmonized descriptions of collections, spatial contexts, and institutional activities. Embedding such semantic models within the federated digital twin architecture facilitates machine-interpretable data exchange, semantic querying, and cross-institutional integration, ultimately enhancing collaboration and long-term data sustainability [23].

Extending the collaborative capability, the **Federation Layer** orchestrates joint operations across multiple institutions. It enables federated learning, coordinated emergency simulations, and adaptive response strategies, fostering strategic alignment among stakeholders such as heritage ministries, emergency services, and international organizations. By maintaining a shared knowledge repository at the federation level, this layer ensures that global insights are effectively tailored to local actions. Anchoring all layers, the **Security Layer** safeguards the architecture by embedding robust data protection mechanisms. These include encryption, access controls, differential privacy, and intrusion detection systems. Its presence is vital in the cultural heritage context, where ethical and legal constraints often govern the sharing of sensitive data, especially regarding sacred artifacts or personal visitor information.

Although initially tested in museums, this framework can be extended to a wide range of cultural heritage contexts. These include archaeological sites vulnerable to environmental degradation, historical urban districts requiring coordinated safety measures, and libraries or archives that house humidity-sensitive manuscripts. It is equally relevant to intangible heritage settings, such as public squares used for cultural festivals or ceremonial gatherings, where real-time crowd monitoring and emergency coordination are essential. By supporting real-time analytics, federated learning, and adaptive feedback mechanisms, the proposed FDT architecture contributes to the ethical, resilient, and intelligent management of both tangible and intangible cultural heritage as presented in Figure 4.2. (See the Appendix for mathematical validation).




Figure 4.2: Multi-layered framework for federated digital twin management in museums.

4.3.1 Simulation-Based Validation

In cultural heritage preservation, particularly in multi-site museum and heritage networks, real-time access to operational or emergency data is often limited due to privacy concerns, logistical constraints, and the infrequency of actual emergencies. Furthermore, conducting real-world evacuation drills or stress-testing coordination mechanisms across multiple institutions poses ethical, financial, and safety challenges, especially when valuable artifacts, historical structures, and visitor well-being are involved [190, 218]. As such, simulation-based validation emerges as a critical and ethically responsible approach for evaluating the proposed Federated Digital Twin (FDT) framework. Simulation enables the controlled modeling of complex emergency scenarios such as fires, earthquakes, or crowd surges within diverse cultural heritage environments without imposing risk on physical sites or individuals. In particular, agent-based modeling (ABM) offers a highly suitable technique for this domain, as it facilitates realistic, granular representation of individual visitor behavior, staff roles, hazard progression, and adaptive evacuation decisions [6, 206]. ABM also allows for the incorporation of contextual factors that are specific to cultural sites, including constrained pathways, preservation zones, varying visitor profiles (e.g., elderly or disabled individuals), and localized emergency protocols. The use of ABM within the simulation framework enables the evaluation of both localized responses and federated coordination strategies across multiple digital twins [164, 218]. It supports experimentation with dynamic parameters such as changing visitor density, evolving hazard intensity, and system-wide communication delays, elements that are difficult to control or replicate in real life. Additionally, simulation permits systematic testing of learning algorithms, feedback loops, and policy interventions within a safe, repeatable environment. This is particularly important for assessing the performance, scalability, and responsiveness of the FDT framework under different stress conditions [57]. For cultural heritage institutions aiming to adopt AI-enhanced decision support systems, simulation provides critical insights into feasibility, interoperability, and operational impact before live deployment [163, 164]. It also serves as a platform for stakeholder engagement, allowing emergency managers, curators, and policymakers to visualize outcomes, test protocols, and refine institutional strategies without compromising heritage assets or public trust. Therefore, the simulation-based validation approach adopted in this study is not only practical and efficient but also aligned with the ethical and operational values of responsible digital heritage management.

Simulation Setup

To empirically evaluate the performance and feasibility of the proposed Federated Digital Twin (FDT) framework within a cultural heritage context, we conducted a series of simulation experiments using NetLogo 6.3.0. NetLogo was selected for its well-established agent-based modeling (ABM) capabilities, which are particularly suitable for representing individual behaviors, social dynamics, and environmental interactions during emergencies [206, 234]. These qualities make it ideal for simulating complex evacuation patterns and decision-making processes that are characteristic of cultural heritage environments.

The simulation environment was designed to emulate a realistic multi-room museum setting populated with 340 virtual visitor agents [194]. Each agent was assigned distinct behavioral attributes, such as mobility speed, reaction time, decision thresholds, and stress levels, based on

empirical findings from crowd behavior studies [92, 266]. The architectural layout of the simulated museum includes exhibition halls, corridors, exits, and preservation-sensitive zones, reflecting the spatial constraints and operational realities typical of heritage institutions. The model focused on evaluating three core components: (1) the behavior of Local Digital Twins (LDTs) in detecting hazards and initiating evacuation decisions, (2) the effectiveness of federated coordination across two distinct museums in sharing learning and optimizing outcomes, and (3) overall evacuation performance under varying emergency conditions, such as fire origin, spread rate, and crowd density. Emergency scenarios were designed to replicate fire incidents, incorporating dynamic elements such as smoke diffusion, reduced visibility, blocked pathways, and agent health degradation due to toxic exposure. These scenarios enabled us to test the responsiveness, adaptability, and collaborative decision-making capabilities of the FDT system. Simulation was employed as a safe, repeatable, and controlled method for assessing the potential impact of AI-enhanced coordination tools, without placing physical heritage assets at risk or disrupting the daily operations of cultural institutions. This approach provides a noninvasive way to explore complex emergency scenarios, test system responsiveness, and refine decision-making strategies. Moreover, the simulation setup lays the groundwork for future integration with 3D digital heritage models and immersive training platforms, thereby supporting ongoing innovation in the protection and management of cultural heritage environments. Simulation experiments were conducted to evaluate the performance of adaptive learning mechanisms within the FDT framework. The scenarios include dynamic fire propagation and varying visitor densities. Key metrics, such as evacuation success rate and route efficiency, were analyzed to assess the impact of adaptive learning on emergency response strategies. To support reproducibility and foster community engagement, the NetLogo simulation model, BehaviorSpace configuration files, and associated analysis scripts are publicly available via GitHub¹ upon publication. This will enable other researchers to replicate our experiments, validate the reported results, and extend the proposed FDT framework for further study.

Simulation Parameters

To ensure that the simulation closely reflected the complexities and operational dynamics of real-world cultural heritage environments, a set of detailed parameters was defined for emergency evacuation scenarios. The virtual environment was modeled as two medium-sized museum facilities populated with 340 agents, representing typical visitors, staff, and support personnel[194]. This population size was selected based on the average visitor volume observed in medium-capacity institutions and was sufficient to produce meaningful congestion patterns and crowd behavior variations. Each agent was designed with individualized attributes, including variable walking speeds, stress-reactive decision-making logic, and interaction thresholds, based on personal space and environmental awareness. These behaviors were informed by empirical research on human behavior during emergencies and calibrated using established parameters from evacuation modeling literature [234]. The architectural layout of the museum includes multiple interconnected rooms, corridors, check points, and designated exit zones. Critical spatial elements, such as visibility, signage availability, bottleneck-prone areas, and physical obstacles (e.g., staircases or artifact displays), were modeled to reflect realistic evacuation challenges. The environment also included dynamic factors, such as diminishing visibility from smoke, changing hazard locations, and mobility impairments,

¹https://anonymous.4open.science/r/FDT_Museum_Evacuation_Simulation-25B5

which collectively influenced both individual agent behavior and overall evacuation dynamics. Multiple emergency scenarios were simulated to evaluate the adaptability and coordination capabilities of the Federated Digital Twin (FDT) architecture. These include randomized fire outbreaks, internal security breaches, and hazard escalation patterns affecting different parts of the museum. During each simulation run, Local Digital Twins (LDTs) operated independently to detect hazards, assess visitor distribution, and initiate evacuation procedures tailored to localized risk conditions. Federated intelligence is simulated through the aggregation of performance metrics and adaptive strategy refinement, allowing simulated learning to occur across runs.

Throughout the simulations, a suite of real-time data was captured for performance evaluation. The key metrics included visitor density over time, average evacuation duration, routing path utilization, and agent health degradation. Special attention was paid to the identification of bottlenecks and decision-point delays, which provided critical insight into spatial limitations, routing algorithm performance, and layout vulnerabilities. This rich dataset forms the basis for a quantitative assessment of the effectiveness of the FDT framework and informs future refinements aimed at improving resilience, responsiveness, and coordination in emergency scenarios within heritage sites.

4.3.2 Empirical Evaluation of the Federated Digital Twin Architecture through Simulation Results

To assess the performance and operational viability of the proposed Federated Digital Twin (FDT) framework for managing emergency evacuations within interconnected museum environments, a comprehensive series of simulation experiments was conducted, as illustrated in Figures 4.3, 4.4, 4.5. These simulations were designed to replicate realistic emergency scenarios under varying conditions while evaluating the key components of the framework: hazard detection, agent response, federated coordination, security enforcement, and adaptive evacuation decision-making. The simulation model was implemented using NetLogo 6.3.0, employing an agent-based modeling approach to represent human behavior during emergencies accurately. The virtual environment was structured to emulate a two-museum configuration, each represented by a Local Digital Twin (LDT). These museums have visitor agents exhibiting diverse behavioral profiles, including speed variability, decision thresholds, and stress sensitivity. Vulnerable groups such as children, elderly individuals, and persons with disabilities were included, each paired with designated helper agents responsible for assisting in evacuation, thereby adding realism and ethical considerations to the model. Although our current simulation involves two museums, the proposed Federated Digital Twin (FDT) architecture is inherently scalable. Its modular design enables additional Local Digital Twins (LDTs) to be integrated seamlessly through the federated learning layer. For larger networks (e.g., five or more institutions), scalability can be maintained using hierarchical federation or clustered coordination, which balance complexity with autonomy and privacy.

In terms of complexity, the integration of new museums increases the computational and communication load in a mostly proportional manner. Each additional site contributes primarily to its own local processing, while the federated layer must coordinate information exchange and synchronization across all institutions. This makes the growth in complexity closer to proportional rather than exponential, especially when hierarchical or clustered coordination strategies are employed. In practice, this ensures that the system remains manageable and efficient even as the number of participating museums grows.

To reflect emergency unpredictability, the hazard configuration was randomized across the simulation runs. Fire origin points were assigned dynamically, and fire propagation followed a tunable rate set at 0.1 per tick for the baseline testing. Smoke behavior was modeled to disperse more rapidly than fire, affecting both visibility and agent health. A linear toxic exposure model quantifies individual health degradation based on proximity to the hazard and duration of exposure, simulating the consequences of delayed response or congestion. Each LDT operates autonomously, leveraging local analytics to detect hazards and trigger evacuation decisions. The decision criteria include visitor density, health indicators, and environmental feedback. Evacuation routing is driven by a shortest-path algorithm with dynamic rerouting capabilities, allowing agents to adjust paths in response to blocked exits, crowd congestion, or evolving hazards. Adaptive behaviors, such as hesitation, panic-induced delays, and group clustering, were also modeled to reflect realistic crowd dynamics during emergencies. The federated component of the simulation was simulated through post-simulation aggregation of the localized outputs. While each LDT processed and executed its operations independently, a layer of federated coordination was simulated by combining the performance metrics across museums. Metrics such as evacuation time, survival rate, response delay, and bottleneck frequency were aggregated to assess system-wide efficiency and collaborative behavior. Shared decision strategies were iteratively refined through scenario replay and used to train a central knowledge base simulating federated learning over time. A security abstraction layer is included in the security and privacy mechanisms of the model. This incorporated access control rules, zone restrictions, and behavioral limits specific to each museum. Simulated intrusion events, such as unauthorized entry into staff-only zones, trigger automated security responses from LDTs. Emergency triggers, including smoke alarms, temperature thresholds, and motion sensors, were embedded within the agents and environmental zones to simulate stimulus-driven responses. Each simulation run was configured to operate with a time step of one tick per second and proceeded until either a complete evacuation occurred or a maximum of 500 ticks was reached. Each scenario was executed 30 times using different random seeds to ensure the statistical robustness and generalizability of the results. Data logging was performed using NetLogo's BehaviorSpace and was extended with custom modules to extract real-time performance data for each architecture layer. Performance evaluation was conducted using a suite of quantitative metrics.

The *evacuation success rate* averaged 93.8% , demonstrating the framework's capacity to guide the vast majority of agents to safety under complex conditions. The *average evacuation time* was measured at approximately 1120 seconds, demonstrating the impact of adaptive routing and coordinated response mechanisms. *Bottleneck occurrences* were consistently recorded around narrow exits or intersecting corridors, providing actionable insights into spatial design limitations and opportunities for architectural improvements in real-world museum layouts. The *overall routing efficiency* was calculated as 88.8%, reflecting the effectiveness of the dynamic rerouting algorithms and congestion management strategies. In our study, routing efficiency was defined as the proportion of agents who successfully followed optimal evacuation routes, obtaining a value of 88. 8%. This indicates that most agents adopted near-optimal routes, while some deviations occurred due to congestion and dynamically evolving hazards. The average evacuation time, computed as the mean time taken by evacuated agents across 30 independent simulation runs incorporating dynamic fire propagation and heterogeneous crowd behavior, was 1120 seconds. These results are consistent with previous evacuation modeling studies. For example, Liu et al.[146] reported a total evacuation time of ap-

proximately 1495 seconds in a large-scale community scenario; Safari et al. [212] observed mean times of around 386 seconds in a primary school setting; and Gerges et al. [82] found average times of about 1056 seconds in a high-rise residential building. Taken together, such benchmarks suggest that evacuation times for medium-sized institutions typically range from 900 to 1500 seconds, placing our findings well within the expected interval as summarized in Table 4.1 .

Table 4.1: Summary of Key Simulation Metrics, Definitions, Benchmarks, and Results for Federated Digital Twin Evaluation

Metric	Definition	Benchmark	Result
Routing Efficiency	% of agents using optimal paths	85–90% typical	88.8%
Avg. Evacuation Time	Mean time to exit	900–1500 seconds	1120 seconds
Evacuation Success Rate	% of agents safely evacuated	90–95% typical	93.8%

Importantly, the Federated Learning aspect of the architecture was modeled using the *Federated Averaging (FedAvg) algorithm*. This enabled efficient model training across distributed digital twins while preserving data locality and reducing communication overhead [231]. By sharing only model parameters rather than raw data, the framework maintained operational privacy and institutional autonomy, which are critical factors in cultural heritage environments in which data ownership and ethical data use are paramount [115]. To ensure the reproducibility and scientific validity of our federated learning setup, we adopted standard configurations aligned with established practices in the literature. Specifically, our choice of learning rate, batch size, local update rounds, and aggregation method follows guidelines outlined in foundational surveys on federated learning [115]. These parameters are optimized for low-client-count environments and have demonstrated stable convergence in simulation-based studies. Table 4.2 summarizes the hyperparameter settings and model architecture used in our experiment.

Table 4.2: Federated Learning Hyperparameters and Model Configuration

Parameter	Value	Description
Model Type	Multilayer Perceptron (MLP)	3 hidden layers (64–32–16 neurons)
Input Features	5	Visitor type, location, route cost, queue length, urgency level
Output	1	Exit decision (route recommendation or nearest exit)
Learning Rate	0.01	Fixed for all clients
Batch Size	32	Used for local training
Local Epochs	5	Training passes per client per round
Global Rounds	20	Total number of communication rounds
Optimizer	SGD	Stochastic Gradient Descent with momentum = 0.9
Aggregation Strategy	FedAvg	Standard weight aggregation approach
Number of Clients	2	Representing the two museums
Data Distribution	Non-IID	Each museum had different visitor flow and density

Overall, the simulation results validated the technical feasibility, operational resilience, and adaptive intelligence of the proposed FDT framework. The experiments demonstrate how federated coordination can improve both local and collective emergency preparedness, while maintaining

the ethical and architectural constraints of sensitive cultural heritage institutions. These findings reinforce the potential of digital twins, when designed with federated intelligence and security layers, to play a pivotal role in the future of smart, resilient, and ethically governed cultural heritage management.




Figure 4.3: Initial agent distribution and hazard zones.




Figure 4.4: Mid-simulation showing evacuation flow.

4.4 Discussion

This study introduced a simulation-based prototype of a Federated Digital Twin (FDT) framework designed to enhance emergency preparedness, real-time coordination, and decentralized decision-making in cultural heritage institutions. The architecture integrates localized digital twins, federated learning, and agent-based modeling to enable intelligent responses across distributed museum environments. The results of our simulation-based validation demonstrate the ability of the system to dynamically guide evacuations with high performance: a 93.8% evacuation success rate and 88.8% routing efficiency. These findings highlight the FDT's potential to optimize evacuation routes, prioritize vulnerable individuals, and improve situational awareness under emergency conditions. To our knowledge, this is one of the first studies to operationalize Federated Digital Twin architectures in the cultural heritage domain, moving beyond theoretical concepts to a practical, simulation-tested framework. The approach addresses critical challenges faced by heritage institutions, including the need to coordinate emergency response across physically distributed sites while




Figure 4.5: Aggregated performance metrics across runs.

preserving data privacy and institutional autonomy. Using federated learning, the system enables real-time collaboration among institutions without requiring centralized access to sensitive data, thus aligning with ethical AI practices and data protection regulations such as the GDPR [151]. A key innovation lies in reimagining digital twins not as isolated monitoring tools but as collaborative, intelligent infrastructures. The modular, scalable architecture is adaptable to diverse heritage contexts, ranging from major national museums and archaeological parks to small, decentralized institutions. The FDT framework can also extend to other domains, such as airports, libraries, or educational campuses, where coordination and privacy are essential.

4.4.1 Case Study Scenario: Coordinated Emergency Response Across Two Museums

To illustrate the real-world applicability of the FDT framework, we present a fictional emergency scenario involving two heritage institutions, Museum A and Museum B, located in the same cultural district but governed independently. At 10:35 a.m., a fire breaks out in the eastern wing of Museum A, home to centuries-old manuscripts. The local digital twin immediately processes sensor data, simulates hazard spread, and initiates evacuation procedures tailored to current visitor flow and agent characteristics. Visitors with limited mobility and school groups are prioritized using the agent profile registry embedded in the local system. The digital twin also updates building maps in real time, rerouting traffic away from blocked exits while ensuring safe and efficient egress.

Simultaneously, without sharing raw data, Museum A transmits encrypted situational metadata to Museum B through the federated digital twin layer. Although not directly affected, Museum B receives the alert and prepares to receive evacuees. It temporarily halts new admissions, deploys staff to standby positions, and updates internal risk models based on real-time federated learning inputs. Throughout this event, the federated infrastructure continuously updates decision models across both sites, maintaining coordination without violating data sovereignty or latency constraints. This scenario mirrors the dynamics simulated in our NetLogo model and demonstrates how a system with a 93.8% simulated evacuation success rate could feasibly operate in a real emergency context.

4.4.2 Limitations and Modeling Realism

Despite its promising potential, the simulation-based validation of the FDT framework comes with inherent limitations. The current agent-based model assumes rational behavior, instant situational awareness, and consistent compliance with evacuation instructions. However, real-world emergencies often involve panic, confusion, and unpredictable behavior that can deviate significantly from logical agent models. Visitors may ignore alarms, help others instead of fleeing, or become immobilized due to fear responses that are culturally and psychologically influenced [267]. From a technical perspective, the model assumes uninterrupted communication between museums, clean sensor data, and perfect building infrastructure. In practice, network interruptions, sensor failures, and misinformation can affect coordination. Furthermore, the simulation's fire propagation model is based on uniform spread parameters, while the actual fire behavior depends on the materials of the building, the ventilation systems, and the vulnerability of the artifact. Health degradation in the model is also linear and lacks the complexity of real physiological responses. These modeling constraints highlight the need to interpret simulation findings as indicative rather than definitive. Continued empirical testing in real environments is essential to validate the framework's scalability and reliability under real-world conditions.

4.4.3 Toward Real-World Deployment

The current prototype corresponds to Technology Readiness Level (TRL) 4, validated through simulation. The next development phase aims to reach TRL 6 through real-world deployment and testing in collaboration with the MuNDA Museum (Museo Nazionale d'Abruzzo ²). This will involve sensor integration, real-time drills, staff training, and performance benchmarking under controlled emergency scenarios. Future work will also include integration of 3D spatial models, Building Information Modeling (BIM), and Extended Reality (XR) interfaces. These immersive technologies will enable spatially accurate simulations, improve user engagement, and support participatory planning involving curators, emergency responders, and visitors. In addition, we aim to align the system with responsible AI principles and FAIR data governance. This includes the development of transparent model auditing, stakeholder-informed coordination protocols, and privacy-preserving model update mechanisms. These efforts will be critical for building institutional trust and fostering long-term adoption across diverse heritage ecosystems. In general, this research contributes a scalable, ethical, and intelligent digital twin architecture that offers strong potential to reshape emergency preparedness and coordination within the cultural heritage sector and beyond.

4.5 Conclusions and Future Work

This study presented a simulation-based prototype of a Federated Digital Twin (FDT) framework designed to enhance emergency preparedness, coordination, and decision-making across distributed cultural heritage institutions. The proposed architecture integrates localized digital twins, federated learning algorithms, and agent-based modeling to enable decentralized, privacy-preserving, and adaptive response strategies in multi-site museum environments. Simulation-based validation using NetLogo demonstrated the system's effectiveness, achieving a 93.8% evacuation success rate and

²<https://museonazionaleabruzzo.cultura.gov.it/>

88.8% routing efficiency under dynamic emergency scenarios. Beyond its technical achievements, the framework contributes to the evolving digital infrastructure of cultural heritage by shifting from passive documentation to proactive, intelligent risk management. The layered and modular design ensures interoperability and scalability, making it applicable to a wide range of tangible and intangible heritage contexts, including museums, archives, archaeological sites, and public cultural spaces. Future development will focus on integrating 3D spatial models, immersive extended reality (XR) technologies, and advanced AI-based threat detection. These enhancements will support immersive training, real-time scenario walkthroughs, and participatory risk assessments, further strengthening institutional resilience and stakeholder engagement. By laying the groundwork for real-world implementation and interdisciplinary collaboration, this research advances the state of the art in digital heritage protection and sets a foundation for ethically governed, intelligent, and sustainable cultural heritage management in an increasingly uncertain world.

Building on the promising results of this simulation-based prototype, several future directions are planned to advance the Federated Digital Twin (FDT) framework toward real-world deployment, interdisciplinary application, and enhanced system intelligence. These efforts will focus on both technical development and operational integration.

1. **Technical Enhancements and AI Integration:** A primary area of future work involves integrating adaptive AI models trained on historical visitor movement and emergency data. These models will enable more precise threat forecasting and personalized evacuation planning based on real-time conditions and user profiles. To improve simulation fidelity, we will also incorporate sensor uncertainty models, probabilistic decision-making algorithms, and more advanced behavioral dynamics. Research will explore questions such as: *How can context-aware AI improve evacuation efficiency across heterogeneous visitor populations?* and *What is the impact of real-time federated learning updates on cross-institutional response times?*
2. **Immersive Interfaces and 3D Digital Twins:** To increase spatial realism and stakeholder engagement, the system will be extended with immersive technologies such as 3D spatial models and extended reality (XR) interfaces. These tools will support emergency walkthroughs, staff training simulations, and participatory planning for curators and first responders. Incorporating Building Information Modeling (BIM) and high-resolution LiDAR scans will enable richer spatial interactions and geospatial accuracy in both training and emergency execution environments.
3. **Real-World Deployment and Empirical Validation:** The current prototype corresponds to Technology Readiness Level (TRL) 4. The next phase aims to advance to TRL 6 through pilot testing in collaboration with the MuNDA Museum and related institutions. This will include field deployment of the FDT modules, real-time evacuation drills, and integration with existing sensor infrastructures. Feedback from museum staff, visitors, and emergency personnel will be collected through human-in-the-loop evaluations and post-event analysis.
4. **Cross-Cultural Extension and Internationalization:** The FDT framework will be adapted and evaluated in varying legal, infrastructural, and cultural contexts. This will help test the scalability and localization of the system in heritage institutions with different resource levels, governance models, and visitor behaviors. Comparative studies across international museum

networks will assess performance consistency and policy compliance in decentralized emergency systems.

5. Ethical Governance, Open Science, and Stakeholder Participation: Future iterations of the system will prioritize ethical AI principles, including transparency, accountability, and alignment with FAIR data standards. A stakeholder-informed governance model will be developed to guide data sharing, federated model training, and response coordination. To support reproducibility and community contribution, simulation models, architectural blueprints, and selected software modules will be released under an open-source license. Overall, these enhancements will guide the evolution of the FDT system from a validated prototype to a robust, deployable platform for intelligent, ethical, and scalable emergency response in cultural heritage environments.

Chapter 5

Real-Time Coordination in Federated Digital Twins

Overview of Chapter

This chapter introduces a novel Real-Time Federated Digital Twin Framework (RT-FDTF) designed to enable deterministic emergency coordination across distributed cultural heritage environments. Building upon the foundational concepts of Digital Twins in Chapter 2, federated learning , and simulation-based validation in Chapter 4, this chapter advances research by integrating Worst-Case Execution Time (WCET) profiling, real-time scheduling, and latency-aware federated synchronization into a unified architecture. The RT-FDTF framework ensures that each local digital twin operates within a strict 500 ms control cycle, executing a six-stage decision loop: sensing, hazard detection, route optimization, synchronization, model update, and broadcasting. Federated learning is implemented using a secure, time-triggered FedAvg protocol, maintaining synchronization latency below 55 ms while preserving data privacy. Key contributions include:

A modular architecture combining embedded control and federated intelligence. Deterministic scheduling using Earliest Deadline First (EDF) and fixed-priority policies. Mathematical modeling of real-time guarantees and synchronization latency bounds. Simulation-based validation using Net-Logo , achieving a 259 ms average cycle time and 89.3% evacuation success rate. The chapter also presents comparative performance analysis, scalability evaluation, and an ablation study, demonstrating the framework’s robustness, efficiency, and applicability to broader industrial information integration scenarios.

5.1 Introduction

The Digital Twin (DT) paradigm has become the cornerstone of next-generation cyber-physical systems, enabling real-time synchronization between virtual models and physical assets [86, 183]. Initially developed for high-fidelity monitoring in aerospace and manufacturing, DTs now serve a broad spectrum of domains, from predictive maintenance in industry to urban planning, energy optimization, and cultural heritage management [9, 78, 202, 240]. As illustrated in Figure 5.1, the concept of Digital Twins (DTs) has evolved significantly over the past five decades, transitioning from NASA’s early aerospace applications to the modern era of Federated Digital Twins (FDTs). This evolution highlights the growing complexity and distributed nature of intelligent systems,

motivating the development of frameworks that support real-time, federated coordination across multiple environments.




Figure 5.1: Evolution of Digital Twin to Federated Digital Twins, showing major technological and conceptual milestones from NASA’s early developments in the 1970s to the expansion of Federated Digital Twin research in 2024.

More recently, the concept of *Federated Digital Twins* (FDTs) has emerged to overcome the scalability, privacy, and interoperability limitations of centralized DT architectures [162, 280]. In an FDT system, multiple localized twins collaborate by sharing abstracted insights such as learned policies or encoded model parameters rather than raw sensor data. This decentralized approach is ideal for environments requiring coordination across organizational or geographic boundaries, including cultural heritage sites, emergency evacuation networks, and distributed industrial assets [229, 244].

Despite this potential, most existing DT and FDT implementations still rely on *soft real-time* behavior, tolerating variable delays without guaranteeing deterministic task completion [3]. Such unpredictability is unacceptable in safety-critical settings where coordination tasks such as fire detection, hazard mapping, or dynamic evacuation routing must complete within fixed deadlines to safeguard human lives and valuable artifacts [156]. Without Worst-Case Execution Time (WCET) analysis or deterministic scheduling, these systems risk cascading latency and potential coordination failure [112].

In cultural heritage environments, such as museums, galleries, or archaeological complexes, emergency events (fires, structural failures, or toxic leaks) demand immediate, cross-site coordination. However, these institutions often operate semi-autonomous digital infrastructures that lack bounded real-time synchronization across distributed sites. This gap creates an *urgent need for architectures that guarantee responsiveness under timing constraints*. To ensure timely and deterministic responses in safety critical cultural environments, this study adopts a 500 ms execution cycle constraint a realistic bound for soft real-time embedded systems where small, bounded delays

are tolerable but timing guarantees remain vital [145, 225].

We therefore propose a novel **Federated Real-Time Digital Twin Integration Framework (FRT-DTIF)** for deterministic emergency coordination in cultural heritage ecosystems. The proposed system integrates task-level WCET modeling, real-time scheduling, secure federated synchronization, and simulation-based validation to ensure each decision loop executes within a guaranteed 500 ms deadline.

Research Questions:

- **RQ1:** How can real-time scheduling and WCET profiling be integrated into federated DT coordination for distributed cultural heritage sites?
- **RQ2:** How effective is the proposed framework in minimizing federated synchronization latency and ensuring bounded coordination among digital twins?
- **RQ3:** Can the proposed architecture be validated through realistic emergency evacuation simulations that reflect multi-site cultural environments?

Contributions:

1. A Federated Digital Twin Integration Framework (FRT-DTIF) that incorporates real-time scheduling and federated coordination for deterministic responsiveness.
2. A six-stage decision pipeline with per-task WCET constraints ensuring a bounded 500 ms decision loop.
3. A secure, time-triggered federated synchronization protocol enabling low-latency metadata exchange between distributed twins.
4. Empirical validation through a NetLogo-based simulation representing federated museum environments and emergency evacuation scenarios.

While the primary use case explored in this study is cultural heritage environments, the proposed FRT-DTIF architecture is equally applicable to industrial domains such as smart manufacturing, logistics hubs, and energy grids. These sectors also require deterministic coordination across distributed cyber-physical systems, making the framework relevant for broader industrial information integration scenarios.

5.2 Background

Digital Twins (DTs) are virtual representations of physical assets that maintain continuous synchronization with real-world processes through bidirectional data exchange [88, 183]. As an integral part of cyber-physical systems (CPS), DTs combine sensor-driven data acquisition, simulation models, and intelligent analytics to enable real-time decision-making and system optimization [9, 202]. They have demonstrated significant value across multiple domains, including smart manufacturing [78], energy systems, intelligent transportation, and cultural heritage preservation, where digital replicas support monitoring of artifacts, buildings, and visitor flows [229, 244].

However, deploying DTs in safety-critical and time-sensitive environments remains challenging. Traditional DT architectures often depend on cloud-centric infrastructures, where unpredictable

communication latency and network congestion can compromise deterministic performance [240]. Moreover, most DT implementations operate as isolated digital entities, lacking mechanisms for secure and real-time collaboration with other twins. These limitations restrict their effectiveness in distributed environments where coordinated responses are required under strict timing and privacy constraints. Consequently, there is a growing demand for decentralized, real-time capable DT frameworks that ensure predictable execution, secure collaboration, and responsiveness under uncertainty particularly in mission-critical domains such as emergency response and cultural heritage protection. Figure 5.1 illustrates this chronological development, showing the major technological and conceptual milestones leading from NASA’s original digital twin idea to contemporary federated digital twin research.

5.2.1 Federated Digital Twins (FDTs)

The Federated Digital Twin (FDT) paradigm extends the DT concept by enabling multiple autonomous twins to collaborate while preserving local data sovereignty and privacy [162, 280]. Rather than centralizing data at a single repository, FDTs exchange abstracted knowledge representations, such as model weights or control policies, using federated learning (FL) mechanisms. This enables cooperative behavior among distributed DTs deployed across geographically separated sites, supporting complex ecosystems such as multi-museum networks or smart cultural heritage infrastructures [57].

Nevertheless, current FDT frameworks primarily emphasize interoperability, model convergence, and data privacy while often overlooking hard or even firm real-time requirements. In time-critical applications such as fire detection, structural monitoring, or evacuation management the presence of non-deterministic network delays and unsynchronized update intervals can lead to degraded system reliability and delayed responses. Therefore, integrating real-time constraints, deterministic scheduling, and timing predictability into federated digital twin coordination is essential to ensure consistent, deadline-bound execution across distributed environments.

5.2.2 Real-Time Systems and WCET Modeling

In real-time systems, correctness depends not only on computational accuracy but also on temporal determinism results must be produced within a specific deadline [83]. Hard real-time systems, such as those used in aerospace, automotive, or industrial robotics, require all tasks to complete before their deadlines, whereas soft real-time systems tolerate bounded delays. To verify schedulability, system designers rely on Worst-Case Execution Time (WCET) analysis, which estimates the maximum possible execution duration of a task under worst-case conditions [156].

Despite its well-established role in embedded control systems, WCET modeling is rarely incorporated into DT or FDT frameworks. Most current implementations depend on cloud or edge orchestration without guaranteeing temporal correctness. Our proposed framework explicitly integrates WCET analysis with real-time scheduling to ensure that all federated DT nodes execute critical coordination tasks such as hazard detection, evacuation routing, and local synchronization within deterministic time bounds. This integration of WCET-aware scheduling and federated synchronization represents a significant step toward enabling dependable, time-bounded, and safe decision-making in distributed cultural heritage environments.

5.3 Related Work

In recent years, extensive research has been conducted on Digital Twin (DT) architectures, federated learning (FL) systems, and distributed coordination models. These advances span multiple domains such as smart cities, industrial automation, energy networks, and edge computing. However, few studies provide a unified framework that combines federated intelligence with formal real-time guarantees and WCET-aware task scheduling, especially for safety-critical cultural heritage applications.

Gil et al. [162] introduced a federated data-sharing architecture designed for urban digital twins, enabling semantic interoperability and decentralized collaboration among city-level infrastructures. While effective in ensuring data privacy and decentralization, their approach lacks mechanisms for deterministic scheduling or deadline compliance, limiting its application in real-time emergency scenarios.

Yang et al. [280] conducted a comprehensive survey of federated machine learning (FL) approaches across heterogeneous distributed systems, emphasizing data heterogeneity, communication efficiency, and model convergence. However, their study remains primarily theoretical and does not address the integration of temporal constraints, bounded inference latency, or coordinated task execution factors essential to real-time federated systems.

Sun et al. [233] proposed an adaptive FL algorithm for industrial CPS, focusing on model adaptability under non-IID data conditions. Although their system improves learning efficiency at the edge, it lacks timing predictability and does not incorporate WCET constraints. Similarly, Zhang et al. [283] explored asynchronous FL for Industrial IoT applications, providing scalability but at the cost of deterministic timing.

Complementary work has explored scheduling and predictability in embedded systems. Kumar et al. [45] applied WCET estimation for control tasks in cyber-physical systems, while Lee et al. [25] designed real-time scheduling mechanisms for embedded devices. These contributions enhance timing reliability but do not address federated intelligence or multi-twin coordination.

Table 5.1 provides a comparative overview of the representative studies, highlighting their support for real-time (RT) guarantees, WCET modeling, and federated learning integration. As shown, most existing works focus on isolated aspects either learning efficiency, system decentralization, or temporal predictability but few address all three simultaneously.

This study fills this gap by introducing a **Federated Real-Time Digital Twin Integration Framework (FRT-DTIF)** that unifies WCET-aware real-time task scheduling, secure federated synchronization, and formal validation through simulation. The proposed framework ensures deterministic, decentralized coordination suitable for distributed cultural heritage environments, bridging the theoretical principles of real-time systems with practical deployment of federated digital twins.

5.4 Federated Real-Time Embedded Digital Twin Integration Framework

The proposed **Federated Real-Time Embedded Digital Twin Framework (FRT-DTIF)**, illustrated in Figure 5.2, introduces a modular and deterministic coordination model for managing emergency responses across distributed cultural heritage environments. This architecture integrates localized real-time control, secure federated synchronization, and latency-aware communication into

Table 5.1: Comparative Analysis of Related Work

Work	Domain	RT Guarantee	WCET Modeling	Federated Learning	Evaluation Approach
Gil et al. (2024) [162]	Urban Digital Twins	×	×	✓	Simulation
Yang et al. (2019) [280]	Federated Learning Theory	×	×	✓	Survey
Sun et al. (2020) [233]	Industrial Edge Systems	×	×	✓	Cloud Simulation
Zhang et al. (2023) [283]	IIoT Federated Twins	Partial	×	✓	IoT Testbed
Lee et al. (2021) [25]	Embedded Real-Time Systems	✓	×	×	Experimental
Kumar et al. (2018) [45]	Scheduling in CPS	×	✓	×	Case Study
This Work	Cultural Heritage DTs	✓	✓	✓	RT Simulation

a unified information integration pipeline. Each component of the framework is designed to operate under a 500 ms decision cycle, enforced through *Worst-Case Execution Time (WCET)* profiling and real-time scheduling analysis.

At each participating cultural heritage site (e.g., a museum or historical building), an autonomous digital twin continuously mirrors the physical environment through a combination of *sensing*, *control logic*, and *actuation* layers. The *sensing layer* collects live data on fire indicators, temperature, humidity, and visitor density. The *control layer* executes deterministic tasks such as hazard detection, route optimization, and decision-making using fixed-priority or Earliest Deadline First (EDF) scheduling. The *actuation layer* enforces decisions through safety mechanisms activating alarms, guiding visitors, or controlling exits. Each of these tasks is profiled with WCET analysis to ensure the entire perception decision action loop completes within the 500 ms deadline, providing predictable system responsiveness even under high load.

Beyond local autonomy, the framework implements a federated coordination layer that enables digital twins to share learned knowledge without transferring raw data. Each site periodically trains a local model on its sensory and decision data, then transmits abstracted updates to a *federated aggregation coordinator*. This coordinator executes a latency-optimized version of the Federated Averaging (*FedAvg*) algorithm [203, 207], combining the local updates into a global model that captures distributed system behavior. The aggregated global model is then disseminated back to all twins, improving their local decision accuracy and adaptation to shared emergency patterns while preserving data privacy.

To ensure timing predictability, the synchronization process operates in a *time-triggered window* of 50–55 ms, synchronized with the real-time execution cycle. This guarantees that model updates do not violate the 500 ms control loop. The scheduler prioritizes synchronization tasks alongside hazard detection and route planning using EDF or fixed-priority rules to maintain both learning continuity and safety guarantees.

A secure communication backbone supports message passing between local twins and the aggregation coordinator. This backbone implements encryption and packet validation to prevent data leakage while maintaining synchronization consistency. The communication middleware tolerates transient network delays and asynchronous updates, ensuring resilience in federated coordination. In addition, an integrated monitoring interface continuously tracks real-time metrics such as synchronization latency, WCET adherence, and evacuation success rates. These indicators enable runtime

verification of system stability and facilitate adaptive reconfiguration if performance degradation is detected. From an information integration perspective, the FRT-DTIF framework enables real-time fusion of distributed knowledge across autonomous digital twins. By abstracting and synchronizing metadata rather than raw data, the system supports semantic interoperability, decentralized decision-making, and scalable coordination key pillars of industrial information integration.

Overall, the FRT-DTIF establishes a deterministic, privacy-preserving, and scalable integration framework that combines real-time embedded control with federated intelligence. This enables distributed cultural heritage institutions to achieve coordinated emergency response and dynamic situational awareness under strict timing and safety constraints.




Figure 5.2: Federated Real-Time Embedded Digital Twin Integration Framework (FRT-DTIF) for Emergency Coordination.

5.4.1 Federated Learning Integration and Coordination

Federated Learning (FL) serves as the primary mechanism for achieving distributed coordination among Digital Twins (DTs) in the proposed framework. Each local DT instance operates autonomously within its physical context (e.g., a museum site) and periodically shares high-level

metadata such as congestion indices, hazard probability, and evacuation efficiency with the global aggregator. By transmitting only abstracted parameters instead of raw sensor streams, the framework reduces bandwidth consumption and ensures compliance with data privacy regulations [139, 223].

The FL process follows a deterministic, time-triggered synchronization protocol. During each control cycle, local models are updated using current observations and then aggregated through the *FedAvg* algorithm [75, 201]. The aggregation window is bounded by a fixed latency threshold (55 ms) to preserve the integrity of the 500 ms decision loop. Once aggregation is completed, the global model is redistributed to local twins, refining their control logic with collective situational awareness.

To maintain hard real-time guarantees, the FL module is tightly integrated with the system scheduler. Synchronization and inference tasks are prioritized using EDF or fixed-priority policies, ensuring temporal determinism. This coupling allows the digital twin network to continuously improve predictive capability while maintaining responsiveness and safety in dynamic, safety-critical environments [27, 144].

5.4.2 Mathematical Model and Real-Time Guarantees

Notation and assumptions. We consider a federation of N Digital Twin (DT) nodes. Each node i executes a local control loop periodically with cycle (control period) T_c (here $T_c = 500$ ms). Within a cycle each node executes m_i periodic tasks:

$$\tau_i = \{(C_{i,1}, T_{i,1}), (C_{i,2}, T_{i,2}), \dots, (C_{i,m_i}, T_{i,m_i})\},$$

where $C_{i,j}$ denotes the (empirical) Worst-Case Execution Time (WCET) and $T_{i,j}$ the period (we assume deadlines equal periods). We further assume:

- Tasks are scheduled preemptively on each node under either EDF or fixed-priority (FP) scheduling.
- A deterministic, time-triggered federated synchronization of duration δ_{sync} occurs once every control cycle T_c and consumes local CPU/network time on each participating node (i.e., it behaves as an additional periodic activity with execution time δ_{sync} and period T_c).
- Network one-way delay is L_{net} and the aggregation/processing time at the coordinator is Δ_{proc} . We ignore secondary queuing effects except where explicitly noted.
- For convergence results we assume standard FL conditions (smooth, possibly strongly convex local objectives, bounded gradients, and appropriate step sizes) and additionally assume bounded staleness induced by δ_{sync} .

Node utilization. For node i define the total (task) utilization

$$U_i = \sum_{j=1}^{m_i} \frac{C_{i,j}}{T_{i,j}}.$$

Including the periodic synchronization overhead δ_{sync} (period T_c) the effective utilization is

$$U_i^{\text{eff}} = U_i + \frac{\delta_{\text{sync}}}{T_c}.$$

[Schedulability under EDF] If each node i schedules tasks under EDF and

$$U_i^{\text{eff}} = U_i + \frac{\delta_{\text{sync}}}{T_c} \leq 1,$$

then all tasks on node i are schedulable (i.e., will meet their deadlines).

[Sketch of proof] Treat the synchronization window as a periodic task of execution time δ_{sync} and period T_c . Under EDF, the necessary and sufficient condition for schedulability of a set of independent preemptable periodic tasks with deadlines equal to periods is that total utilization does not exceed 1 (Liu & Layland). Thus if $U_i^{\text{eff}} \leq 1$ the task set is schedulable.

Response-time analysis under fixed priority. Under fixed-priority preemptive scheduling a standard response-time recurrence for a task $\tau_{i,j}$ is:

$$R_{i,j}^{(k+1)} = C_{i,j} + \sum_{\tau_{i,h} \in hp(\tau_{i,j})} \left\lceil \frac{R_{i,j}^{(k)}}{T_{i,h}} \right\rceil C_{i,h} + I_{\text{sync}}, \quad (5.1)$$

where $hp(\tau_{i,j})$ is the set of higher priority tasks and $I_{\text{sync}} = \delta_{\text{sync}}$ models the interference due to the synchronization window (treated as a periodic higher-priority job). Iterate (5.1) from $R_{i,j}^{(0)} = C_{i,j}$ until convergence; $\tau_{i,j}$ is schedulable if the fixed point $R_{i,j} \leq D_{i,j}$.

Synchronization latency bound. A more accurate aggregation latency (one synchronization round) can be lower-bounded by network and processing delays. A conservative bound for a single federation round (local upload $\rightarrow$ aggregation $\rightarrow$ global broadcast) is

$$L_{\text{sync}} \leq 2L_{\text{net}} + \Delta_{\text{proc}} + L_{\text{ser}},$$

where L_{ser} accounts for serialization/queuing effects at the aggregator (which may scale with N for naive implementations). If the aggregation node processes updates in parallel and the network supports concurrent transfers, L_{ser} is small; in serialized aggregation L_{ser} grows approximately linearly in N (or with message size). In our design we assume a low-latency aggregator (parallel processing) so that

$$L_{\text{sync}} \approx 2L_{\text{net}} + \Delta_{\text{proc}}.$$

Thus to satisfy synchronization budget δ_{sync} we require $L_{\text{sync}} \leq \delta_{\text{sync}}$.

Convergence comment for federated averaging. Under standard assumptions for FedAvg (smoothness, bounded variance of stochastic gradients, bounded delays/staleness, and appropriate learning rates), the aggregated model converges to a neighborhood of the centralized solution; the convergence rate depends on problem smoothness and heterogeneity [33, 203]. In the presence of bounded synchronization delay δ_{sync} (treated as bounded staleness), FedAvg retains convergence guarantees with a degradation factor proportional to the staleness and data heterogeneity. A precise

convergence rate requires problem-specific assumptions; hence we state the following qualitative corollary.

[Qualitative stability of FedAvg under bounded delay] If every node is schedulable (Theorem 5.4.2) and synchronization delay satisfies $L_{\text{sync}} \leq \delta_{\text{sync}}$ with δ_{sync} small relative to T_c , then FedAvg with appropriate step size achieves convergence similar to the synchronous case up to an error term that depends on model heterogeneity and delay. Detailed quantitative bounds follow from standard analyses of FedAvg under bounded staleness [151, 203].

Design rule (practical). For our target parameters $T_c = 500$ ms and the measured $\delta_{\text{sync}} \leq 55$ ms, a sufficient design rule is

$$U_i \leq 1 - \frac{\delta_{\text{sync}}}{T_c} \approx 0.89,$$

i.e., allocate at most $\sim 89\%$ utilization to local tasks. If observed utilization approaches this bound, apply overload-mitigation (reduce task complexity, defer noncritical computation, or increase cycle period).

5.5 Simulation Setup and Evaluation

Given the complexity, safety sensitivity, and decentralized nature of the proposed Federated Real-Time Digital Twin Integration Framework (FRT-DTIF), simulation-based validation was adopted as the most practical and ethically responsible approach for evaluating real-time performance, coordination behavior, and system scalability under emergency conditions.

5.5.1 Rationale for Simulation-Based Evaluation

Simulation is a widely recognized methodology for validating cyber-physical and embedded systems, particularly when real-world experimentation in safety critical domains such as healthcare, transportation, or emergency evacuation is infeasible or ethically constrained [14, 235]. Within cultural heritage environments, conducting live fire drills or hazard experiments poses unacceptable risks to visitors, staff, and irreplaceable artifacts. Simulation therefore provides a controlled, reproducible, and scalable environment to analyze real-time coordination and timing compliance without endangering physical assets.

Moreover, simulation facilitates systematic exploration of variables such as visitor density, hazard propagation speed, and digital twin configuration. This allows for parametric sensitivity analysis and stress testing of real-time scheduling under varying workloads. Isern et al. [106] highlight that simulation-based analysis is especially effective for studying multi-agent social behaviors, adaptive decision-making, and emergent coordination capabilities essential for validating federated digital twin systems.

Formally, each simulation cycle was monitored to record the following time dependent metrics:

$$\text{Cycle Delay} :: t_{\text{end}} - t_{\text{start}}, \quad (5.2)$$

$$\text{WCET_Violation} :: \begin{cases} 1, & \text{if Cycle Delay} > 500 \text{ ms,} \\ 0, & \text{otherwise.} \end{cases} \quad (5.3)$$

These indicators, along with synchronization latency and throughput, enabled quantitative verification of timing constraints established in Section 5.4.2. Such measurable metrics are critical in certifying compliance with real-time requirements before physical deployment [127].

Choice of NetLogo as Simulation Platform

NetLogo was selected as the simulation platform due to its robustness in modeling multi-agent systems (MAS) and its proven effectiveness in replicating decentralized decision-making processes [205, 271]. The agent-based paradigm directly maps onto the digital twin architecture, where each museum, visitor, or hazard entity operates autonomously but remains coupled through federated synchronization. This structure allowed the accurate emulation of distributed digital twins communicating via time-triggered coordination.

NetLogo provides built-in primitives for spatial reasoning, agent mobility, and neighbor-based interactions, enabling the modeling of complex phenomena such as panic propagation, visitor clustering, and evacuation path adaptation. These capabilities, combined with BehaviorSpace batch experimentation and integration with Python through extensions, facilitated large-scale parametric studies of timing and coordination performance. Similar approaches have been demonstrated in federated simulation frameworks for smart city and digital heritage applications [201].

Real-time data visualization in NetLogo enabled both qualitative and quantitative validation of system behavior. Using the platform’s built-in metric export and plotting functions, we monitored:

- Average and maximum decision cycle times per twin;
- Frequency and duration of WCET violations;
- Synchronization latency between federated nodes;
- Aggregate evacuation success rates and congestion resolution times.

These outputs provided a comprehensive dataset for assessing performance reliability, timing consistency, and coordination scalability across two interconnected virtual museums. NetLogo’s reproducibility features and transparent visual interface also supported interdisciplinary collaboration, aligning with best practices in digital humanities and smart heritage research [246].

Taken together, the combination of simulation methodology and the NetLogo platform provided a realistic, ethically sound, and information-integrative testbed for validating the FRT-DTIF. It allowed stress testing of real-time scheduling policies, verification of federated learning latency, and assessment of emergency coordination effectiveness under diverse cultural heritage scenarios.

5.5.2 Real-Time Evaluation and WCET Profiling

To evaluate the real-time performance of the proposed Federated Real-Time Digital Twin Integration Framework (FRT-DTIF) in safety-critical cultural heritage environments, an extensive simulation-based timing analysis was performed. The primary focus was on three interdependent metrics: (1) control cycle duration, (2) Worst-Case Execution Time (WCET) compliance, and (3) federated synchronization latency. Together, these parameters define the framework’s ability to integrate decentralized decision-making processes under strict temporal constraints.

5.5.3 Task-Level Execution and Cycle Profiling

Each Digital Twin (DT) instance, representing an autonomous museum node, operates within a 500 ms real-time control cycle. This cycle is composed of six sequentially dependent stages: (1) sensor acquisition, (2) hazard detection, (3) evacuation route planning, (4) federated synchronization, (5) model update, and (6) instruction broadcasting. These stages collectively form the operational pipeline for real-time decision-making.

Empirical profiling using NetLogo’s runtime instrumentation revealed the average execution time per task: 45 ms (sensor acquisition), 80 ms (hazard detection), 120 ms (routing), 50 ms (federated synchronization), 90 ms (model update), and 30 ms (instruction broadcast), yielding a cumulative mean of 415 ms per cycle. This leaves an operational margin of approximately 85 ms sufficient to absorb moderate jitter or task overruns within the 500 ms real-time constraint, consistent with soft real-time system requirements [112, 225].

The use of dual scheduling strategies Fixed Priority (FP) and Earliest Deadline First (EDF) ensured that task deadlines were respected under normal load conditions. Across 1120 simulation ticks, the system maintained an average control cycle duration of 259 ms, while 119 cycles exceeded the WCET limit, corresponding to an 11.9% violation rate (Table 5.2). These violations occurred primarily under high-density evacuation phases when route planning computations became more complex, indicating areas where adaptive scheduling or load distribution mechanisms could further enhance determinism.

Table 5.2: Cycle Performance Summary (1120-Tick Simulation)

Metric	Observed Value
WCET Threshold	500 ms
Mean Cycle Time	259 ms
Synchronization Latency	22 ms
Total WCET Violations	119
Violation Frequency	11.9%

WCET Profiling Methodology

The WCET analysis followed a semi-formal profiling approach using the `timer` primitive in NetLogo to measure execution durations with millisecond granularity. Profiling was conducted under varying visitor densities and hazard intensities to capture both nominal and stress conditions. For each task, the maximum observed runtime across all ticks was recorded as an empirical approximation of its WCET. While this method lacks static code-level analysis, it provides a realistic bound suitable for soft real-time system validation [272].

The system defines a cycle violation function:

$$V(t) = \begin{cases} 1, & \text{if } D_t > 500 \text{ ms,} \\ 0, & \text{otherwise,} \end{cases}$$

where D_t represents the measured duration of cycle t . The total number of violations V_{tot} and the

average cycle time $\bar{D}$ were computed as:

$$V_{tot} = \sum_{t=1}^T V(t), \quad \bar{D} = \frac{1}{T} \sum_{t=1}^T D_t.$$

This formulation enables a quantitative link between empirical profiling and the schedulability guarantees derived in Theorem 1 (Section 5.4.2). Observing $\bar{D} < 500$ ms validates that system utilization remains within the theoretical bound $U_i + \frac{\delta_{sync}}{T_c} \leq 1$.




Figure 5.3: Real-time Federated Digital Twin Control Loop and WCET Profiling.

5.6 Simulation Outcomes and Performance Analysis

The full-scale simulation involved two virtual museums represented as independent DTs, each hosting a population of 170 agents including visitors, staff, and vulnerable individuals. All nodes operated in synchronous 500ms control cycles coordinated through a federated synchronization layer.

Cycle Timing and WCET Compliance: The average execution cycle was 259ms, with peak WCET measured at 634ms during severe congestion. A total of 119 WCET violations were recorded (11.9%), corresponding to moments of route recomputation and hazard diffusion. Despite these transients, more than 88% of cycles completed within deadlines, demonstrating bounded response under dynamic load. This behavior aligns with soft real-time operational tolerances and validates the feasibility of bounded-latency coordination in safety-critical environments.

Federated Synchronization Latency: The average synchronization delay between the two DT nodes was 22ms, substantially below the 55ms ceiling established in Lemma 1. This confirms that the federated averaging protocol maintained real-time compatibility without degrading local cycle execution. Latency stability across cycles also suggests robustness against short-term processing spikes, ensuring consistent integration across federated nodes.

Evacuation Efficiency and Safety Metrics: Across 1120 ticks, the system achieved an 89.3% evacuation success rate, representing agents successfully reaching safe zones within simulated deadlines. Dynamic route recalculation and adaptive congestion management improved throughput by approximately 18% compared to static routing baselines. This performance surpasses prior models such as Zhang et al. [283], which lacked dynamic adaptability and timing awareness. Yang et al. [280] emphasized federated learning efficiency but did not address deadline-constrained evacuation. Hence, the FRT-DTIF framework introduces a unique integration of federated coordination and real-time decision control.

Visualization and Monitoring: Figure 5.4 illustrates real-time monitoring outputs, including WCET violations, synchronization delays, and evacuation trends. The visualization underscores the temporal predictability and cross-twin consistency of the system, highlighting its applicability to cultural heritage emergency coordination.




Figure 5.4: Simulation dashboard showing WCET cycle time, synchronization latency, and evacuation performance.

5.6.1 Extended Results and Comparative Evaluation

To complement the baseline timing and evacuation metrics, this section provides a deeper analysis of the proposed Federated Real-Time Digital Twin Integration Framework (FRT-DTIF) from three

perspectives: (1) comparative performance against benchmark systems, (2) scalability under increasing federation size, and (3) ablation analysis to quantify the contribution of each architectural component. The objective is to establish empirical evidence of efficiency, robustness, and scalability core evaluation dimensions for industrial information integration systems [280].

5.6.2 Comparative Performance Analysis

The proposed FRT-DTIF was compared against three baseline systems to assess its benefits in timing, synchronization, and evacuation outcomes:

- **Centralized DT:** A single digital twin performing all computations and coordination, representing traditional monolithic control architectures.
- **Asynchronous Federated DT (Async-FDT):** A distributed model without timing alignment, where synchronization occurs at variable intervals.
- **Non-WCET-Aware FDT:** A federated system without real-time scheduling or WCET enforcement, relying solely on best-effort task execution.

Table 5.3 summarizes the comparative results averaged over 10 independent runs with different random seeds. The proposed FRT-DTIF achieved the lowest average decision latency and highest evacuation success rate, confirming that bounded WCET control and federated learning complement each other for both temporal and operational efficiency.

Table 5.3: Comparative performance of FRT-DTIF versus baseline systems (mean $\pm$ 95% CI)

System	Mean Cycle Time (ms)	Sync Latency (ms)	Evacuation Success (%)	WCET Violation (%)
FRT-DTIF (proposed)	259 $\pm$ 12	22 $\pm$ 3	89.3 $\pm$ 2.1	11.9
Centralized DT baseline	392 $\pm$ 20	80 $\pm$ 8	74.5 $\pm$ 3.8	28.4
Async-FDT baseline	310 $\pm$ 15	48 $\pm$ 6	82.0 $\pm$ 3.0	18.7
Non-WCET-Aware FDT	285 $\pm$ 14	24 $\pm$ 4	86.1 $\pm$ 2.5	15.3

Statistical analysis using one-way ANOVA confirmed significant differences ($p < 0.05$) between the FRT-DTIF and all baselines in mean cycle time and evacuation rate. These results emphasize the integrated benefits of deterministic scheduling and federated model adaptation, both of which are critical for dependable industrial and cultural heritage applications [34, 94].

5.6.3 Scalability Evaluation under Increasing Federation Size

To assess scalability, we simulated additional nodes representing digital twins for multiple heritage sites, gradually increasing the federation size from $N = 2$ to $N = 10$. Each node operated under identical timing constraints and exchanged model updates every five cycles.

Synchronization latency exhibited near-linear growth with federation size but remained below the theoretical 55 ms threshold for up to $N = 8$ nodes. Beyond this, minor deviations emerged due to communication contention. The mean control cycle duration increased modestly from 259 ms (for $N = 2$) to 296 ms (for $N = 10$), confirming that the system scales gracefully without violating real-time guarantees.

This behavior is consistent with analytical expectations, since total utilization $U_i + \frac{\delta_{sync}}{T_c}$ remained below unity for all configurations. The scalability results affirm that the federated synchronization

mechanism can accommodate multisite coordination without centralized bottlenecks, satisfying the integration scalability criterion.

5.6.4 Ablation Study on System Components

To quantify the contribution of individual modules, an ablation study was performed by selectively disabling the core components, namely WCET scheduling and federated learning synchronization, while keeping all other configurations constant. As summarized in Table 5.4, the absence of either mechanism resulted in substantial degradation of timing stability and evacuation outcomes.

Table 5.4: Ablation study of system components

Configuration	Cycle Time (ms)	Evacuation (%)	WCET Violations (%)
Full FRT-DTIF	259	89.3	11.9
Without WCET scheduling	305	84.1	19.7
Without Federated Sync	292	81.2	17.3
Without Both	370	73.8	28.6

The results demonstrate that WCET enforcement contributes primarily to temporal stability (reducing violation frequency), while federated synchronization enhances operational effectiveness by promoting collective intelligence. Together, these mechanisms deliver synergistic performance improvements across all metrics.

5.6.5 Interpretation and Industrial Relevance

The extended evaluation confirms that the FRT-DTIF architecture not only satisfies real-time constraints but also scales across multiple federated entities with minimal performance degradation. These findings directly support the core vision of industrial information integration enabling interoperable, time-bounded collaboration among autonomous digital twins distributed across organizations and geographic locations [33, 144, 232]. The modularity of the framework allows seamless integration with existing industrial standards such as OPC UA and IEC 61499, facilitating interoperability with legacy systems [136]. This compatibility ensures that the FRT-DTIF can be adopted incrementally within existing industrial infrastructures without requiring complete system overhauls. From a deployment perspective, the linear scalability trend and bounded synchronization overhead imply that cultural heritage networks or smart city clusters can deploy federated coordination layers without the need for high-bandwidth infrastructure. Edge-level execution of 500 ms cycles and 20–25 ms synchronization periods are achievable using off the shelf embedded devices, making the approach feasible for real-world deployment in low-resource environments.

Future extensions of this analysis will evaluate performance under non-ideal networking conditions, heterogeneous compute resources, and partial federation scenarios, offering further validation of the framework’s resilience and practical adaptability.

5.7 Discussion and Implications

The results presented in this study demonstrate the feasibility and practical relevance of deploying a Real-Time Federated Digital Twin (FRT-DTIF) framework for coordinated emergency management

in distributed environments. Beyond its immediate cultural heritage use case, the architecture introduces fundamental design principles applicable to a wide range of industrial information integration scenarios.

5.7.1 Theoretical and Architectural Implications

The proposed FRT-DTIF framework provides a new perspective on how real-time systems and federated intelligence can converge within cyber-physical environments. By enforcing bounded execution times through WCET profiling and real-time scheduling, the system bridges a long-standing gap between theoretical determinism and practical distributed coordination. This combination ensures predictable latency across federated nodes an aspect often overlooked in existing federated learning or digital twin implementations [201, 286].

Furthermore, the inclusion of federated learning within an embedded real-time loop advances the concept of "self-learning twins" that evolve through continuous feedback without compromising timing constraints. This aligns with emerging research trends in distributed edge intelligence, where devices must simultaneously guarantee responsiveness and autonomy [33, 144]. The FRT-DTIF framework thus contributes to a unified architecture that is both time-aware and knowledge-adaptive.

5.7.2 Practical and Industrial Implications

Although evaluated within the domain of cultural heritage, the design principles underpinning this framework generalize naturally to industrial information integration systems. In smart factories, for example, multiple production lines can act as federated digital twins, collaboratively optimizing safety, efficiency, and energy consumption without centralizing sensitive data. The same real-time coordination strategies can support predictive maintenance, adaptive scheduling, or collaborative robotics each requiring low-latency decision-making under distributed constraints.

The results also demonstrate that decentralized control combined with federated synchronization significantly reduces coordination latency compared to centralized architectures. This property is vital for industrial networks that span geographically distributed plants, where network reliability, autonomy, and security are crucial [33]. The system's ability to maintain synchronization latency under 55 ms and execution cycles below 500 ms provides strong evidence of its readiness for edge-level industrial deployment.

5.7.3 Integration with Industrial Information Systems

From an information integration perspective, the FRT-DTIF framework represents a scalable middleware layer that supports both vertical and horizontal integration. Vertically, it enables interoperability between embedded controllers, digital twins, and enterprise-level management systems. Horizontally, it allows distributed twin instances each representing local assets or processes to collaborate via federated learning while preserving data privacy and autonomy. Such integration patterns are increasingly central to industrial standards like RAMI 4.0 and ISO 23247, which define interoperability between digital twins and production systems.

Furthermore, the proposed architecture aligns with the industrial movement toward "system-of-systems" coordination, where multiple digital twins exchange meta-information rather than raw

data. This approach reduces bandwidth requirements and enhances cybersecurity, addressing common challenges in large-scale Industry 4.0 deployments [219].

5.7.4 Societal and Sustainability Implications

Beyond industrial applications, the FRT-DTIF framework holds relevance for societal resilience and sustainability. In the context of cultural heritage preservation, the approach offers a pathway to protect human life and irreplaceable assets through predictive emergency coordination. More broadly, federated digital twins can facilitate energy-efficient resource management, reduce waste in industrial production, and support the monitoring of sustainable development indicators, aligning with the objectives of SDGs 9 (Industry, Innovation, and Infrastructure) and 11 (Sustainable Cities and Communities).

5.7.5 Limitations and Future Directions

While simulation-based validation provided strong empirical insights, several limitations must be acknowledged. The current implementation does not account for network-induced jitter, packet loss, or heterogeneous computational resources. Moreover, human behavioral complexity such as panic propagation or compliance variability was simplified in the agent models. These limitations may affect scalability in real-world deployments.

Future research should therefore explore hybrid validation approaches that combine simulation, hardware-in-the-loop (HIL) testing, and pilot deployments within operational environments. Additional work is also required to formalize the system's real-time guarantees through analytical verification methods such as response time analysis (RTA) or model checking. Integrating more advanced federated optimization algorithms such as FedProx, SCAFFOLD, or adaptive asynchronous FL will further improve robustness under non-IID data distributions and heterogeneous latencies.

5.7.6 Broader Research Implications

The findings extend the theoretical frontier of digital twin research by linking federated learning and real-time embedded control. They demonstrate that synchronization, decision-making, and learning can coexist under strict timing constraints, a property critical to future industrial information systems that must balance autonomy, scalability, and reliability. Consequently, this research provides a blueprint for real-time distributed learning in industrial ecosystems, bridging the gap between academic theory and practical implementation.

5.8 Conclusions and Future Work

This study proposed and validated a comprehensive Real-Time Federated Digital Twin Integration Framework (FRT-DTIF) designed for time-critical coordination and emergency management in cultural heritage environments. The framework represents a significant advancement in industrial information integration, combining real-time embedded control, federated intelligence, and agent-based behavioral modeling into a cohesive architecture capable of distributed decision-making under strict timing constraints.

The key contribution of this research lies in its holistic integration of real-time scheduling, Worst-Case Execution Time (WCET) profiling, and federated synchronization within a unified digital twin framework. By establishing a 500 ms control cycle and embedding deterministic timing guarantees, the system ensures predictable coordination among autonomous museum twins. The simulation-based evaluation demonstrated that the architecture maintained an average execution cycle of 259 ms, with synchronization latency averaging 22 ms well below the 55 ms theoretical bound derived in the analytical model. These quantitative findings affirm the feasibility of achieving low-latency, deadline-compliant coordination across distributed heritage sites.

Furthermore, the integration of federated learning into the real-time loop enabled privacy-preserving collaboration among multiple digital twins without centralizing sensitive data. The system achieved an evacuation success rate of 89.3%, with adaptive behaviors driven by helper agents and predictive route selection improving throughput by nearly 18% compared to static evacuation models. These outcomes highlight the potential of combining distributed learning with real-time control to enhance safety and resilience in critical infrastructure.

From an information integration perspective, the proposed framework operationalizes multi-source data fusion and decision synchronization across independent entities, supporting the broader vision of interoperable, adaptive digital ecosystems. It provides a reference model for integrating federated analytics with embedded decision cycles, directly contributing to the research domains of real-time digital twin coordination, edge-based intelligence, and cyber-physical system (CPS) integration.

Despite these promising outcomes, several limitations remain that will guide future research. First, the current simulation assumes idealized communication conditions; future work will incorporate realistic network dynamics including latency variability, packet loss, and bandwidth constraints. This will be achieved through hybrid co-simulation with network-level tools such as ns-3 or OM-NeT++. Second, the agent models can be further enriched using deep reinforcement learning (DRL) or adaptive pathfinding algorithms, allowing the system to autonomously learn evacuation strategies from historical behavior and continuous feedback.

Third, the current framework focuses on horizontally distributed coordination between two museum environments. Extending this to multi-building or multi-floor topologies with vertical mobility (stairs, elevators) will enable scalability to complex architectural layouts such as airports, archaeological sites, and interconnected galleries. Additionally, the integration of privacy-preserving techniques such as differential privacy and secure aggregation will ensure compliance with emerging data protection standards in federated digital ecosystems.

Finally, a prototype implementation on embedded edge computing platforms (e.g., Raspberry Pi or NVIDIA Jetson) will be pursued in collaboration with real-world cultural heritage institutions. This hardware-in-the-loop (HIL) validation will allow empirical verification of WCET guarantees, synchronization stability, and evacuation performance using live sensor streams and occupancy data. Such a deployment will bridge the gap between simulation and field readiness, establishing a concrete pathway for the adoption of real-time federated digital twins in operational museum safety systems.

In conclusion, the FRT-DTIF framework contributes a foundational approach to real-time information integration, merging theoretical rigor, distributed intelligence, and practical deployability. By ensuring deterministic coordination across federated digital twins, it provides a scalable and

privacy-aware model for safeguarding both human life and cultural assets. Beyond cultural heritage, the proposed FRT-DTIF framework presents a reusable blueprint for industrial information integration systems requiring real-time coordination. Future work will extend the architecture to smart manufacturing, transportation, and energy sectors, where federated real-time learning can improve safety, reliability, and resilience. Collaborative field deployments with partner institutions will also be pursued to validate timing constraints, privacy guarantees, and interoperability in real-world industrial settings.

Chapter 6

Cognitive Federated Digital Twins for Sustainable Cultural Heritage

Overview of Chapter

This chapter presents the most advanced contribution of the dissertation: the development and validation of a Cognitive Federated Digital Twin (CFDT) framework designed to support sustainable, privacy-preserving and intelligent cultural heritage management. Building upon the foundational concepts of Digital Twins Chapter 2, AI-driven Federated Digital Twin 4, and real-time coordination Chapter 5, this chapter introduces a self-learning, adaptive architecture that integrates Federated Reinforcement Learning, the MAPE-K autonomic control loop, and multi-agent simulation. The CFDT framework enables distributed museums to operate autonomous digital twins powered by Proximal Policy Optimization (PPO) agents. These agents learn locally and collaborate globally via FedProx-enhanced federated learning, ensuring robust convergence under non-IID data conditions. The embedded MAPE-K loop dynamically adjusts learning rates, action scaling, and strategy preferences based on real-time feedback from operational indicators such as energy efficiency, artifact preservation, visitor comfort, and safety. Key contributions include: a scalable CFDT architecture combining cognitive intelligence, federated learning, and self-adaptive control. Integration of the MAPE-K loop for dynamic optimization across heterogeneous museum environments. Simulation-based validation using a hybrid NetLogo-Python environment across four virtual museums (A-D), demonstrating: 38% improvement in energy efficiency, 23% increase in preservation quality, 18% faster gradient convergence, Coordination efficiency exceeding 3700% and Low communication overhead (0.08–0.35 MB per round)

This chapter incorporates, with editorial adaptation, the core content of the paper “A Cognitive Federated Digital Twin Framework Empowered by Self-Learning Agents for Sustainable Cultural Heritage Preservation,” submitted to the ICT4S 2026 Conference, as well as the ICSA 2026 conference paper “Architecting Cognitive Federated Digital Twins with LLM Agents for Privacy-Preserving Adaptive Systems.” Figures, tables, and simulation results have been reformatted to align with the dissertation’s structure while preserving the integrity of the original scientific contributions. References have been harmonized with the thesis-wide bibliography, and cross-references have been added to ensure coherence with other chapters.

6.1 Introduction

Cultural heritage institutions such as museums play a vital role in preserving and disseminating collective human history. However, maintaining optimal environmental conditions, ensuring visitor safety, and promoting sustainability across multiple heritage sites present complex operational challenges [101]. The emergence of *Digital Twin* (DT) technology provides a promising avenue to address these issues by creating a virtual replica of a physical environment capable of continuous monitoring, simulation, and prediction [86, 129].

Despite these advances, traditional digital twin systems often operate in isolation, limiting their ability to learn collaboratively across different sites. Furthermore, concerns regarding data privacy, heterogeneity of sensing infrastructure, and limited computational resources hinder centralized learning approaches. To overcome these challenges, *Federated Learning* (FL) has recently emerged as a distributed machine learning paradigm that enables multiple agents or institutions to train shared models without exchanging raw data [139, 172]. When combined with digital twins, federated learning provides a mechanism for cross-institutional collaboration that preserves privacy and enhances knowledge sharing [203, 286].

At the same time, self-adaptive systems in the Internet of Things (IoT) and cyber-physical environments are increasingly adopting the *Monitor-Analyze-Plan-Execute over a shared Knowledge* (MAPE-K) model to enable autonomous decision-making and dynamic optimization [75]. Integrating the MAPE-K loop into federated digital twin systems allows real-time adaptation of operational parameters such as energy consumption, visitor comfort, and artifact preservation quality [218].

Building on these advances, this research proposes a *Cognitive Federated Digital Twin (CFDT)* framework that synergistically integrates federated reinforcement learning with the MAPE-K loop for multi-museum collaboration. The CFDT model enables distributed self-learning among several museum digital twins (labeled Museum A–D in this study), allowing them to cooperatively enhance environmental management, preservation integrity, and energy efficiency.

In this research, each museum operates as an autonomous cognitive agent that monitors its environment (temperature, humidity, vibration, and visitor flow) and learns adaptive control strategies through local reinforcement learning. A federated aggregation mechanism ensures global knowledge convergence while maintaining data privacy. The embedded MAPE-K layer dynamically tunes learning parameters such as local learning rate and control scaling based on continuous feedback from performance metrics like comfort, preservation stability, and safety.

The proposed CFDT framework introduces a scalable cognitive federated architecture designed to support sustainable, privacy-preserving, and intelligent heritage management across networked cultural sites. In this study, four virtual museums collaborate through federated learning to autonomously optimize key operational indicators, including visitor flow, energy efficiency, preservation conditions, and safety performance. To ensure privacy and confidentiality of sensitive museum sensor data, the framework integrates secure aggregation and federated learning principles, reducing data exposure risks while preserving local autonomy.

Beyond conceptual design, this work employs a hybrid simulation environment that combines NetLogo for agent-based museum dynamics with a Python Gymnasium-based PPO-FedProx learning engine. This hybrid approach enables real-time interaction between autonomous agents and reinforcement-driven decision intelligence, allowing each museum to adapt its control strategies col-

laboratively without sharing raw state information. By aligning federated digital twin intelligence with privacy-aware learning mechanisms and the MAPE-K loop, the CFDT architecture establishes a foundation for secure, adaptive, and data-driven cultural heritage operations in smart city ecosystems.

6.1.1 Research Questions

The present study investigates how cognitive and federated intelligence can enhance the sustainability and operational resilience of distributed cultural heritage institutions. The research is guided by the following key questions:

1. **RQ1:** How can cognitive reinforcement learning mechanisms, when embedded in a Federated Digital Twin (FDT) architecture, improve real-time decision-making for energy efficiency, comfort, and preservation in multi-museum environments?
2. **RQ2:** In what ways can the integration of the MAPE-K feedback loop enable self-adaptive optimization across heterogeneous museum conditions without centralized data sharing?
3. **RQ3:** How does the inclusion of FedProx regularization affect model convergence, communication efficiency, and coordination across distributed digital twins operating under non-IID data conditions?

These questions aim to bridge the current gap between adaptive learning theory, digital twin engineering, and sustainable cultural heritage management through an integrated cognitive architecture.

6.1.2 Contributions

This work presents a holistic framework that unites cognitive control, federated reinforcement learning, and digital twin modeling for sustainable heritage management. The main contributions are as follows:

1. **A novel Cognitive Federated Digital Twin (CFDT) architecture:** Each museum operates as an intelligent digital twin with self-learning PPO agents that interact with local environmental parameters (visitor density, energy usage, preservation risk, safety index) and collectively optimize global sustainability performance through federated coordination.
2. **Integration of MAPE-K for adaptive self-optimization:** The MAPE-K continuously analyzes system behavior, generates adaptive plans, and autonomously tunes local learning rates and control scales, enabling cognitive adaptability to dynamic environmental or operational changes.
3. **FedProx-enhanced federated learning:** The use of FedProx regularization mitigates client heterogeneity and gradient drift during global aggregation, ensuring stable convergence of distributed models under varying museum conditions and data non-uniformity.

4. **Sustainability-driven multi-agent simulation:** A multi-museum simulation was implemented using Python’s `Gymnasium` and PyTorch frameworks, representing four federated heritage sites (Museums A–D). Each twin adapts its energy, preservation, and safety policies via cognitive reinforcement learning.

6.2 Background

The evolution of Digital Twin (DT) technology has transformed how physical environments are simulated, monitored, and optimized by enabling dynamic virtual replicas that continuously synchronize with real-world systems through sensor data and computational models [66, 203]. While traditional DTs primarily support real-time visualization and predictive analytics, their reliance on centralized data pipelines and static control logic limits scalability and adaptability in distributed and privacy-sensitive domains such as cultural heritage institutions. Federated Digital Twins (FDTs) extend this paradigm by enabling multiple decentralized DT instances to collaboratively learn global patterns without sharing raw data, thereby supporting confidentiality, model generalization, and cross-institutional decision-making across geographically dispersed environments. However, FDTs still struggle to autonomously adapt to shifting environmental conditions, unexpected disruptions, and nonlinear preservation-comfort-energy dynamics. Cognitive Digital Twins (CDTs) address these challenges by incorporating cognitive intelligence mechanisms such as reinforcement learning, semantic reasoning, and autonomic feedback loops that empower digital replicas to sense contextual changes, analyze performance, plan adaptive strategies, and execute corrective actions while continually expanding a shared knowledge base [6, 57]. The convergence of DT, FDT, and CDT technologies forms a powerful foundation for future heritage preservation systems that are capable of sustainable self-optimization, privacy-preserving collaboration, and resilient cross-site intelligence, making them well-suited to the operational complexities of multi-museum ecosystems managing diverse environmental, visitor flow, and conservation conditions [164].

6.2.1 Digital Twin Technology in Cultural Heritage

The concept of the *Digital Twin* (DT) has evolved as a cornerstone of Industry 4.0, bridging the physical and virtual worlds through continuous synchronization and data exchange [46]. In cultural heritage contexts, DTs have been employed to monitor artifact conditions, simulate environmental dynamics, and enhance visitor engagement through virtual and augmented reality [275]. By integrating sensor networks, 3D modeling, and data analytics, digital twins enable curators to analyze and predict the effects of environmental parameters such as temperature, humidity, and light on sensitive artworks [170].

However, most heritage-focused DTs remain static or reactive, lacking the cognitive ability to adapt autonomously to evolving operational states [86]. This limitation constrains their effectiveness in large-scale, multi-site cultural ecosystems where local conditions vary significantly. For instance, interconnected museums with different microclimates and visitor densities require dynamic adaptation mechanisms and cross-site learning to maintain optimal environmental balance.

6.2.2 Federated Learning for Distributed Knowledge Sharing

Federated Learning (FL) introduces a decentralized paradigm where multiple agents train shared machine learning models collaboratively without sharing raw data [75, 172, 280]. This approach addresses privacy and governance issues often encountered in cultural institutions managing sensitive environmental and visitor data. In the museum context, FL allows each site (e.g., Museum A–D) to train local models using its own sensor data while contributing to a globally aggregated model maintained by a central server or coordinating agent.

Federated approaches have shown promise across smart manufacturing [27], healthcare [207], and intelligent transportation [286] domains. However, their application in cultural heritage management remains underexplored. Integrating FL with DTs offers unique potential for scalable collaboration across geographically distributed heritage sites, enabling adaptive optimization of resource usage, preservation conditions, and visitor experience.

6.2.3 Cognitive and Self-Learning Digital Twins

The evolution from traditional DTs to *Cognitive Digital Twins* (CDTs) introduces self-awareness and autonomous learning capabilities that enable continuous system improvement [157, 218]. Cognitive twins leverage reinforcement learning (RL) and semantic reasoning to dynamically adjust control strategies in response to changing environmental and operational conditions. These systems can perceive states, analyze performance metrics, plan adaptive actions, and execute control adjustments while continuously updating their knowledge base [218].

Recent studies have demonstrated the integration of artificial intelligence and IoT within DT frameworks to enable predictive maintenance, anomaly detection, and adaptive control in manufacturing and infrastructure systems [66]. Extending these capabilities to the museum domain allows for self-learning models that can autonomously balance artifact preservation requirements with energy efficiency and visitor comfort.

6.2.4 The MAPE-K Loop for Self-Adaptation

The *Monitor-Analyze-Plan-Execute over Knowledge* (MAPE-K) framework, introduced by IBM, serves as a foundational model for autonomic and self-adaptive systems [164]. The MAPE-K loop enables continuous monitoring of system states, analysis of contextual information, planning of corrective actions, and execution of adjustments informed by a shared knowledge repository. When integrated into digital twin architectures, MAPE-K enhances feedback-driven decision-making and promotes resilience in dynamic environments [66].

In the proposed research, the MAPE-K loop is extended into a distributed cognitive layer that operates at both local (per museum) and global (federated) levels. Each museum agent continuously monitors local sensor data, analyzes deviations, plans adaptive strategies through reinforcement learning, and executes actions to restore equilibrium. The federated coordinator aggregates these learned policies into a global cognitive model, updating the shared knowledge base for system-wide optimization.

6.2.5 Research Gap and Contribution

While digital twin and federated learning technologies have been individually explored, their convergence into a unified cognitive and self-learning framework for multi-site cultural heritage management remains limited. Existing works often focus either on centralized simulation or on isolated predictive analytics without inter-institutional collaboration or autonomy.

This study addresses these gaps by proposing a *Cognitive Federated Digital Twin (CFDT)* framework for distributed museum ecosystems. The CFDT combines federated reinforcement learning, the MAPE-K adaptation loop, and self-learning agents to achieve privacy-preserving collaboration, dynamic adaptation, and continual performance improvement. Through a simulation involving four virtual museums (Museum A–D), this work demonstrates how cognitive federated mechanisms can enhance operational efficiency, safety, and sustainability in the digital preservation of cultural heritage.

6.3 Methodology

6.3.1 Overview of the Cognitive Federated Digital Twin (CFDT) Framework

The proposed *Cognitive Federated Digital Twin (CFDT)* framework integrates Digital Twin (DT), Federated Learning (FL), and the MAPE-K (Monitor-Analyze-Plan-Execute over Knowledge) loop to enable adaptive and collaborative intelligence among multiple museum environments. Each museum site (A–D) operates an autonomous Local DT, equipped with sensing, simulation, and a reinforcement learning (RL) decision layer. Local models learn optimal control policies based on visitor density, energy usage, artifact preservation requirements, and safety conditions.

Rather than sharing raw sensor data, each site participates in privacy-preserving collaborative training via federated learning. The global coordinator aggregates model updates to produce a shared cognitive knowledge repository and redistributes improved policies to local sites. This enables learning synergy across culturally heterogeneous environments while ensuring security, confidentiality, and operational autonomy.

6.3.2 System Architecture and Components

The CFDT system comprises five primary layers:

1. **Physical Layer:** Real museum spaces with sensors for climate, energy consumption, visitor occupancy, air quality, microclimate stress indicators, and emergency signals.
2. **Digital Twin Layer:** A dynamic virtual replica of each museum that simulates environmental behavior, thermal responses, crowd motion, and artifact preservation sensitivity.
3. **Cognitive MAPE-K Layer:** The adaptive intelligence core, executing the MAPE-K feedback loop to enable self-monitoring, anomaly detection, predictive planning, and optimal actuation. A semantic knowledge base stores pattern insights, artifact vulnerability ontology, and cross-museum adaptation heuristics. The cognitive layer distinguishes CFDT from traditional digital twins by enabling reasoning, learning, and self-adaptation based on accumulated operational knowledge.




Figure 6.1: Proposed Cognitive Federated Digital Twin (CFDT) architecture for interconnected museums. Local DTs interact with physical layers and reinforcement learning agents, while the cognitive MAPE-K loop enables adaptation. A secure federated layer aggregates knowledge without exposing private data.

4. **Sustainability Optimization Layer:** A meta-cognitive module that integrates energy efficiency, carbon-awareness, and preservation-centered objectives into the decision-making pipeline. It evaluates energy consumption patterns, microclimate stability, artifact stress indicators, and visitor comfort states to generate multi-objective sustainability plans. A dedicated Sustainability Knowledge Base stores low-energy strategies, carbon-aware policies, and cross-museum sustainability insights, allowing the system to continuously refine environmentally responsible operational behaviors.
5. **Federated Learning Coordination Layer:** Secure, distributed learning across museums. Global model updates enforce collaborative intelligence while guaranteeing local data privacy. Knowledge is shared via encrypted model weights and performance logs.

6.3.3 MAPE-K Driven Adaptive Reinforcement Learning Loop

Each museum agent continuously executes the MAPE-K loop:

- **Monitor:** Collect real-time environmental variables (temperature, humidity, energy usage, density).
- **Analyze:** Detect deviations, identify risks, estimate future stress loads.
- **Plan:** Select optimal control targets (airflow, cooling effort, crowd routing).
- **Execute:** Apply actions via DT actuators or real HVAC/crowd control systems.
- **Knowledge:** Store patterns, decisions, and improvement outcomes.

The policy is defined as:

$$a_t = \pi_\theta(s_t) = \arg \max_a Q_\theta(s_t, a)$$

$$\theta_{t+1} = \theta_t + \alpha \nabla_\theta \mathbb{E}[R(s, a)]$$

In the above formulation, s_t denotes the state of the museum environment at time t , including visitor density, energy level, preservation risk, and safety indicators. The agent selects an action a_t according to a policy $\pi_\theta(s_t)$, parameterized by θ , where the policy represents the probability distribution over possible control actions such as operational modes and environmental adjustments. The function $Q_\theta(s_t, a)$ represents the expected long-term reward of taking action a in state s_t under policy parameters θ . The learning rule updates the policy parameters by moving in the direction of the gradient $\nabla_\theta \mathbb{E}[R(s, a)]$ scaled by the learning rate α . Intuitively, α regulates how aggressively the model adapts to new information, while θ_{t+1} reflects the improved policy after one optimization step. Together, these components allow the agent to iteratively refine decisions, balancing comfort, energy efficiency, preservation, and safety objectives unique to dynamic museum environments.

6.3.4 Federated Reinforcement Learning with FedProx

Each museum trains its local PPO-based RL model for multiple epochs. After local optimization, model weights are sent to the global coordinator. Unlike classical Q-learning or deterministic control

rules, PPO enables stable policy optimization under dynamic and uncertain museum conditions by constraining policy updates and preventing catastrophic shifts in behavior. To address non-IID environmental dynamics and heterogeneity across museums (different crowd behavior, building structures, artifacts, climate sensitivities), we adopt the FedProx objective [139]:

$$\min_w F_i(w) + \frac{\mu}{2} \|w - w_t\|^2$$

where w_t is the global model and μ regulates deviation from the global policy. This stabilizes learning under different museum profiles. We evaluate $\mu \in \{0.001, 0.01, 0.1\}$.

The global update follows:

$$w_{t+1} = \sum_{i=1}^N \frac{n_i}{n_{total}} w_t^{(i)}$$

6.3.5 Reward Function

$$R_t = 0.35C_t + 0.25E_t + 0.20P_t + 0.20S_t$$

Priorities: comfort (C), energy efficiency (E), preservation (P), safety (S).

6.3.6 Reinforcement Learning and PPO Policy Optimization

The RL backbone of CFDT is based on Proximal Policy Optimization (PPO), a stable on-policy actor-critic method suited for continuous and discrete action control. Each museum maintains an actor-critic architecture, where the actor predicts optimal climate and visitor-flow control actions, while the critic evaluates the value of current environmental states.

PPO optimizes a clipped surrogate objective to prevent excessively large policy updates, ensuring stable learning in non-stationary museum environments. This is crucial, as environmental changes such as sudden visitor influx or micro-climate fluctuations could lead to unstable updates if standard policy gradient approaches were used.

The control actions proposed by the policy include: (1) a discrete operational mode (conservation, balanced, performance), and (2) continuous regulation parameters for climate and crowd control.

These multi-dimensional action spaces reflect real-world museum operational constraints, where decisions involve balancing artifact safety, visitor comfort, and energy policies. The combination of PPO and FedProx yields a resilient and decentralized reinforcement learning mechanism suited to cultural heritage management.

6.3.7 Digital Twin Simulation and Adaptation Logic

The Digital Twin serves as a high-fidelity virtual replica of the museum physical environment, simulating thermal conditions, human movement patterns, and artifact vulnerability dynamics. Each museum's DT generates synthetic operational episodes that emulate crowd behavior, environmental transitions, and risk patterns, enabling safe experimentation before control policies affect the real system.

During each local simulation cycle, environmental states evolve according to the selected actions and natural disturbances. The agent observes visitor density, control energy usage, humidity and temperature deviations, artifact stress variables, and safety markers. The DT provides immediate feedback on the consequences of actions, supporting learning without risking cultural assets.

By embedding MAPE-K inside the RL loop, the digital twin does not merely predict outcomes it autonomously analyzes operational deviations and plans adaptive corrections. This transcends traditional DTs by integrating reasoning, learning, and self-optimization, resulting in a cognitive DT capable of progressive intelligence growth.

6.3.8 Baseline Methods

To ensure scientific validity, the CFDT system is compared against two baseline configurations:

(1) Local-Only Reinforcement Learning (No FL): Each museum trains a local PPO agent without knowledge sharing. This baseline measures the benefit of cross-site coordination.

(2) Standard FedAvg: Federated PPO without FedProx regularization. This baseline evaluates robustness against non-IID visitor and environmental conditions.

These configurations isolate the contributions of (i) federated RL, (ii) FedProx, and (iii) the cognitive MAPE-K loop.

6.3.9 Simulation Environment and Execution Protocol

To validate the CFDT in a controlled yet realistic setting, we employed a hybrid co-simulation strategy where NetLogo models the micro-level museum dynamics (visitors, room layout, crowd movement, and emergency behavior), while Python (Gymnasium + PyTorch) controls cognitive decision-making through PPO and FedProx-based federated learning. NetLogo and Python communicate at each simulation step through NetLogoLink, enabling AI-driven adaptive policies to directly influence agent behavior in the virtual museums. The simulation models four independent museums, each represented as a cognitive digital twin with unique micro-climate constraints and visitor movement dynamics. To emulate realistic operational diversity, each museum is initialized with distinct thermal coefficients, humidity sensitivity parameters, and visitor flow distributions.

Each digital twin executes reinforcement learning episodes independently, where stochastic visitor arrivals, motion patterns, and environmental fluctuations introduce dynamic uncertainty. Environmental noise is sampled from a Gaussian distribution $N(0, \sigma^2)$, where σ varies between 0.01 and 0.05 to simulate fluctuating conditions [106, 201]. The maximum number of steps per episode is set to 60, representing one simulated operational period.

Training proceeds for 10 global federated communication rounds. In each round, every digital twin performs 6 local PPO updates before exchanging model parameters with the federated server. Hyperparameters are summarized in Table 6.1. The learning rate is adaptively scaled using the MAPE-K planning module, and FedProx regularization $\mu \in \{0.001, 0.005, 0.01\}$ is applied to mitigate client drift.

All experiments were executed on a system with an AMD Ryzen 7 CPU and NVIDIA RTX GPU, with an average simulation time of 12.4 minutes per full run. To ensure robustness, each experiment was repeated three times with different random seeds, and mean values are reported. The simulation infrastructure ensures that no real visitor or artifact data were used; all interactions

occur within the digital twin environment.

Table 6.1: Hyperparameters for CFDT Simulation

Parameter	Value
Environment Steps / Episode	60
Federated Rounds	10
Local PPO Epochs / Round	5
Discount Factor γ	0.99
Clip Ratio (PPO)	0.2
Learning Rate Range	$1e^{-4} - 2e^{-4}$ (adaptive)
FedProx Regularization μ	$\{0.001, 0.005, 0.01\}$
Exploration Noise	$N(0, \sigma^2)$, $\sigma \in [0.01, 0.05]$
Optimizer	Adam
Action Space	Mixed discrete + continuous
Hardware	AMD Ryzen 7, NVIDIA RTX GPU
Repeats / Seeds	3 runs

Table 6.2: Metric definitions and computation method

Metric	Definition
Energy Efficiency	$EE = 1 - \frac{\text{energy_used}}{\text{energy_capacity}}$ averaged across museums and rounds
Preservation Quality	$PQ = 1 - \text{preservation_risk}$, where risk accumulates with visitor density and humidity deviations
Communication Overhead	Size of serialized model weights exchanged per round in MB
Coordination Efficiency Index	$\frac{\text{mean global reward}}{1 + \text{mean comm. cost}} \times 100$
Gradient Convergence	Relative reduction in loss norm across FL rounds

6.3.10 Algorithm CFDT Training Process

[1] Initialize global model w_0 each round $t = 1 \dots T$ each museum i in parallel Receive global weights w_t local epochs Run PPO step, compute local loss $F_i(w)$ Add proximal term $\frac{\mu}{2} \|w - w_t\|^2$ Update weights $w_t^{(i)}$ Update MAPE-K knowledge memory Send $w_t^{(i)}$ to server Aggregate: $w_{t+1} = \sum \frac{n_i}{n_{total}} w_t^{(i)}$ Broadcast w_{t+1} to all museums

6.4 Findings

The experimental evaluation of the CFDT framework yielded several significant findings that demonstrate its technical robustness and sustainability potential. It achieved consistent performance improvements across four digital museum environments.

6.4.1 Quantitative findings

Energy Efficiency: Improved by 36–39% over local-only RL, measured as the normalized reduction in energy consumption relative to capacity limits.

Preservation Quality: Increased by 21–24% due to adaptive humidity-visitor policy control from MAPE-K intervention.

Comfort & Safety Stability: Comfort index maintained within 0.73–0.84 and safety index within 0.91–0.97 across sites.

Communication Efficiency: Communication load ranged between 0.08–0.35 MB per round, demonstrating lightweight FL exchange suitable for bandwidth-constrained cultural institutions.

Coordination Efficiency: The coordination index stabilized inside the 3520–3860% range, indicating improved decision consistency despite heterogeneous museum conditions.

Convergence Behavior: A mean loss reduction of 17–19% was observed, confirming stable federated PPO adaptation.

6.4.2 Qualitative Findings

Cognitive Adaptability: The MAPE-K layer successfully adjusted local learning rates and exploration parameters in response to environmental drift, enabling each museum twin to autonomously recover from suboptimal states.

Collaborative Intelligence: The federated learning setup enabled each digital twin to benefit from collective knowledge without exposing local data, aligning with GDPR-compliant cultural data management practices.

Resilience and Sustainability: The integration of adaptive control and federated aggregation established a resilient digital ecosystem that balances preservation needs, visitor comfort, and energy efficiency in real time. **Scalability of Framework:** The CFDT’s modular design and lightweight communication protocol make it adaptable to additional heritage sites or sensor networks, supporting large-scale smart cultural ecosystems.

Interpretation

Overall, the findings confirm that the CFDT framework provides a robust and scalable foundation for sustainable digital twin ecosystems. By merging reinforcement learning with federated cognition, the system not only improves performance and sustainability metrics but also introduces a self-adaptive intelligence layer that aligns technological innovation with cultural and environmental preservation objectives.

6.5 Results and Discussion

6.5.1 Overview of Simulation Outcomes

The experimental simulation was executed using the Cognitive Federated Digital Twin (CFDT) framework across four virtual museums, namely Museum A, Museum B, Museum C, and Museum D. Each museum acted as an autonomous digital twin, equipped with localized environmental data streams including temperature, humidity, vibration, and visitor flow. The simulation ran over ten federated communication rounds, with each round comprising five local reinforcement learning epochs. The objective was to collaboratively optimize four core operational parameters: *comfort*, *energy efficiency*, *preservation*, and *safety*.

The output of the simulation demonstrated progressive adaptation across all participating museums, as reflected by increasing average rewards per round and converging MAPE-K metrics. In the early rounds, fluctuations were observed in both reward and control stability due to initial exploration phases of the reinforcement learning agents. However, as the federated cycles progressed, the CFDT system exhibited notable convergence, characterized by a consistent increase in both reward stability and energy optimization.

Table 6.3 summarizes the average local rewards per museum per round, along with the aggregated MAPE-K analysis that highlights the evolving balance among comfort, energy efficiency, preservation, and safety.

Local-Only vs Federated Results

To further validate the effectiveness of the CFDT framework, we conducted an ablation experiment comparing four scenarios across Museums A–D: (i) local-only PPO training without collaboration, (ii) standard FedAvg-based federated RL, (iii) FedProx-enhanced federated RL, and (iv) a beeline (baseline) fixed-policy controller that applies static climate comfort parameters without cognitive adaptation. The results demonstrate that CFDT significantly outperforms isolated learning and non-proximal FL strategies, achieving faster convergence, higher cumulative reward, and improved stability across heterogeneous museum environments. In contrast, the beeline policy exhibited limited adaptability, resulting in lower preservation and energy-efficiency outcomes, confirming the value of cognitive reinforcement learning. Figure 6.2 illustrates the average episodic reward progression, where CFDT shows a steeper learning trajectory, reduced oscillation, and superior equilibrium performance compared to the baselines.

6.5.2 Convergence and Learning Behavior Analysis

The overall performance trajectory demonstrates a steady improvement in the average federated reward, signifying effective policy refinement over successive communication rounds. Initially, the reward oscillations indicated exploration dominance in local policy search, typical of reinforcement learning agents adapting to stochastic museum environments. After the third federated round, the global model began to converge, implying that the shared model parameters effectively captured cross-site patterns of environmental dynamics and visitor behavior.

Figure 6.3 illustrates the federated reward evolution across rounds. The curve follows an upward trend, stabilizing around the seventh round. This stabilization indicates the transition from




Figure 6.2: Performance comparison between Local-Only RL (Museums A-D) and the proposed CFDT framework. The CFDT system achieves higher cumulative rewards and smoother convergence, demonstrating improved learning stability and cross-museum knowledge transfer.

Table 6.3: Summary of Federated Learning Results across Ten Rounds

Round	A	B	C	D	Comfort	Energy	Preservation	Safety
1	40.95	41.58	41.75	41.45	0.65	0.32	0.97	0.80
2	42.42	39.77	44.50	43.15	0.64	0.37	0.96	0.85
3	41.84	43.14	43.25	42.06	0.62	0.46	0.94	0.85
4	43.32	41.28	42.17	43.62	0.62	0.53	0.93	0.84
5	44.28	45.19	42.48	44.43	0.64	0.54	0.94	0.83
6	44.71	45.16	46.41	42.62	0.66	0.58	0.94	0.84
7	42.68	47.08	44.58	40.36	0.65	0.64	0.92	0.81
8	46.68	45.23	41.83	43.54	0.64	0.72	0.88	0.80
9	42.76	44.94	42.27	43.50	0.64	0.76	0.84	0.80
10	47.34	48.03	45.70	42.84	0.66	0.85	0.82	0.77

exploratory learning to exploitation of learned control strategies, reflecting a mature adaptation stage of the CFDT.

Notably, the average reward increased from approximately 41.4 in the first round to 45.97 by the tenth round, demonstrating a relative performance gain of about 11%. This improvement is attributed to the effective integration of the MAPE-K adaptive mechanism, which dynamically tuned learning rates and action policies in response to the system’s observed state changes.

6.5.3 Role of the MAPE-K Loop in System Adaptation

The MAPE-K loop acted as a cognitive meta controller, enabling self-awareness and continuous recalibration of operational parameters. In particular, the *Analyze* and *Plan* stages were crucial in adjusting the optimization strategies between rounds. After each federated aggregation, the




Figure 6.3: Federated reward convergence across 10 rounds for the CFDT system.

MAPE-K plan refined the local learning rates, as observed in the applied plans:

$$\text{opt}_{i,\text{lr}} \in [0.00016, 0.00010]$$

This dynamic reduction of learning rates corresponds to the stabilization phase of the global model, a well-established practice in reinforcement learning to improve convergence speed and avoid overfitting [2, 214].

As rounds progressed, the MAPE-K analysis also showed a consistent upward trend in the energy efficiency metric from 0.32 in round 1 to 0.85 in round 10 reflecting more optimal energy management behaviors. Simultaneously, the comfort and preservation scores maintained stable high values (around 0.65–0.95), validating the system’s ability to balance multiple objectives without degrading artifact safety.

The safety metric, while slightly fluctuating, maintained an average of 0.82 across rounds. This outcome suggests that the system successfully mitigated unsafe conditions (e.g., visitor congestion) through adaptive policy enforcement while maintaining preservation priorities.

6.5.4 Inter-Museum Collaborative Learning Dynamics

A distinctive feature of the CFDT approach is its federated learning mechanism, which facilitates collective intelligence without centralizing sensitive museum data. During each communication round, local models learned environment-specific control policies, while the global aggregation synthesized cross-domain insights. This knowledge fusion enhanced generalization across diverse museum settings, reflecting the robustness of federated reinforcement learning architectures [139, 280].

Museums A and B consistently exhibited higher learning gains throughout the simulation, with their average rewards surpassing 47 by the tenth round. This trend indicates that these environments probably had richer state variations or better conditioned reinforcement dynamics, enabling faster convergence. Museums C and D, although slightly slower, benefited indirectly from the global model synchronization, achieving noticeable improvements by later rounds.

This cooperative interaction can be interpreted as a form of *federated transfer learning*, where localized experience of one digital twin enhances the learning efficiency of others. Such collaboration mirrors real-world heritage management, where operational insights from one museum (e.g., optimal humidity control or visitor flow regulation) can inform better decisions at others without explicit data sharing. Figure 6.4 illustrates the coordination efficiency achieved across ten federated training rounds. The efficiency remained consistently above 3500%, reflecting strong collaboration and knowledge transfer between decentralized museum digital twins. Notably, efficiency peaked in rounds 8–10 as the PPO+FedProx policy stabilized, confirming that the MAPE-K adaptation contributed to improved global learning stability.




Figure 6.4: Coordination efficiency trends across federated rounds. The CFDT framework maintains high coordination efficiency ($\approx 3550\% - 3850\%$) demonstrating stable collaborative learning across distributed museum environments.

6.5.5 Analysis of Communication Efficiency and Latency

Communication efficiency is a critical consideration in federated architectures, especially when scaling to large heritage networks. The simulation reported average communication volumes ranging from 0.08 MB to 0.35 MB per round with latency between 0.0013s and 0.0049s. These values confirm that the CFDT implementation is computationally lightweight and network-efficient, making it practical for real-world deployment where bandwidth or infrastructure may be constrained. Figure 6.5 depicts the relationship between communication volume and overall system efficiency. It is observed that even with marginal increases in data exchange (e.g., round 5 and 9), the global model achieved higher synchronization accuracy without significant latency overhead.

The efficiency ratio, defined as reward improvement per communication cost, averaged 3680%, signifying that each communication round resulted in multiple-fold performance gains relative to the transmitted model updates. This is particularly beneficial for resource-constrained cultural heritage environments, where edge devices may have limited computational and network capacity.




Figure 6.5: Communication volume vs. system efficiency across federated rounds.

6.5.6 Behavior of Multi-Objective Optimization

The CFDT system optimizes multiple competing objectives comfort, energy efficiency, preservation, and safety. These objectives often exhibit trade-offs; for example, increasing visitor comfort (via temperature control) can increase energy consumption, while maximizing artifact preservation might constrain comfort levels. The MAPE-K-driven cognitive adaptation effectively balanced these trade-offs by dynamically re weighting the multi-objective reward function:

$$R_t = \beta_1 C_t + \beta_2 E_t + \beta_3 P_t + \beta_4 S_t$$

In early rounds, the planner prioritized preservation and safety (with higher β_3 and β_4), ensuring artifact integrity. As the system matured, weights gradually shifted toward energy optimization and visitor comfort, reflecting the transition to a steady-state operational regime.

This dynamic balance demonstrates the cognitive capability of the CFDT framework, aligning with prior research emphasizing self-adaptive control in digital twins [46, 139]. The preservation score's marginal decline (from 0.97 to 0.82) reflects this planned rebalancing rather than a performance degradation.

6.5.7 Comparative Discussion with Prior Research

Compared to existing digital twin implementations for cultural heritage [201, 207], the CFDT framework introduces three key advancements:

1. **Federated Reinforcement Learning Integration:** Unlike traditional centralized models, CFDT allows decentralized learning while achieving near-global optimality. This is critical for museums that operate independently under strict data governance laws.
2. **Cognitive Self-Adaptation via MAPE-K:** The integration of MAPE-K grants the system self-awareness and continuous adaptation, addressing limitations of static control systems in

dynamic museum environments.

3. **Multi-Objective Sustainability Optimization:** The reinforcement learning policy simultaneously considers visitor comfort, energy, safety, and artifact preservation, aligning with SDG 7 and sustainable cultural heritage preservation goals.

Furthermore, prior federated learning applications in smart cities or industry [144, 203] typically focus on homogeneous environments (e.g., factories or vehicles). The CFDT demonstrates the feasibility of applying these principles to heterogeneous, mission-critical cultural heritage systems where environmental, human, and preservation dynamics interact in complex ways.

6.5.8 Interpretation of MAPE-K Evolution and Cognitive Maturity

By analyzing the MAPE-K output metrics, it becomes evident that the system’s cognitive maturity increased progressively across the simulation. Comfort and preservation maintained high levels, while energy efficiency nearly tripled from round 1 to round 10. This trajectory signifies that the system learned to allocate resources optimally, improving sustainability without compromising artifact safety.

The observed safety fluctuations can be interpreted as adaptive risk-balancing responses. When the system experimented with new control policies (e.g., adjusting visitor flow thresholds), safety temporarily dipped but recovered as the policy stabilized. Such behavior aligns with the principles of self-learning cyber-physical systems described by [67, 172, 223], confirming that cognitive control can tolerate short-term deviations for long-term gains.

6.5.9 Scalability and Generalization Implications

The scalability of the CFDT approach lies in its federated architecture. Since local agents operate asynchronously, the system can accommodate additional museums or cultural sites without structural reconfiguration. The linear communication cost and low latency make it suitable for large-scale heritage networks.

Moreover, generalization across museums implies that the learned global policy can transfer to new institutions, accelerating deployment and reducing training time. For instance, if a new museum (Museum E) with similar environmental conditions joins the network, it can initialize its local twin using the globally aggregated model, thereby inheriting prior knowledge accumulated from Museums A–D.

This federated transfer mechanism has substantial implications for regional and continental digital heritage preservation initiatives, particularly in Europe, where museums operate under varying infrastructural constraints but share common conservation goals.

6.5.10 Resilience and Stability of the Federated Digital Twin System

Resilience analysis revealed that the CFDT maintained stable operation even under simulated communication delays and partial agent failures. In test scenarios where one museum (e.g., Museum C) temporarily missed a communication round, the global model continued to converge, and local recovery occurred within two subsequent rounds. This resilience is a direct result of the redundancy and autonomy inherent in the federated structure.

Furthermore, the model's energy efficiency improvement from 0.32 to 0.85 across ten rounds underscores its robustness against energy constraints. The preservation index remaining above 0.80 throughout the simulation confirms that the system safeguards critical assets under varying environmental perturbations.

6.5.11 Interpretation of Final Knowledge State

By the tenth federated round, the CFDT achieved a balanced equilibrium across all objectives, with comfort (0.66), energy efficiency (0.85), preservation (0.82), and safety (0.77). These values reflect a mature, optimized cognitive system capable of maintaining operational excellence under complex multi-agent interactions.

The final MAPE-K state suggests the system achieved self-regulation and anticipatory control, meaning it could predict and mitigate potential disruptions before they occurred. This characteristic mirrors cognitive computing theories in autonomic systems, reinforcing that digital twins can transcend static mirroring toward proactive management. Federated reinforcement learning was integrated with the NetLogo-based multi-museum digital twin to validate adaptive decision-making across four cultural heritage environments. Each museum instance learned locally to optimize visitor density, energy usage, cultural preservation conditions, and safety levels, while periodically contributing model updates to a global federated model without sharing raw data. As shown in Figure 6.6, the cumulative reward curves for Museums A–D exhibit stable upward trends with mild fluctuations, demonstrating the agent's ability to adapt to real-time conditions and heterogeneous museum contexts. The learning stability and convergence confirm that the CFDT improves operational balance across dimensions, achieving high reward levels (72–82 range) across all museums. These results validate the feasibility of applying federated reinforcement learning to multi-site cultural heritage facilities, where privacy-preserving collaboration enhances decision quality in dynamic visitor environments. Figure 6.7 shows the spatial configuration of the four simulated museums (A–D) used in the cognitive federated digital twin environment. Each museum consists of three exhibition rooms with distinct visitor flow and environmental dynamics. The rooms are color-coded for clarity (Museum A: red, Museum B: blue, Museum C: green, Museum D: yellow). This layout reflects realistic heterogeneous museum spatial structures and supports controlled experimentation for visitor movement, stress propagation, and environmental management.

6.5.12 Summary of Insights and Implications

The experimental findings validate that the Cognitive Federated Digital Twin framework effectively enables distributed learning, cognitive adaptation, and sustainable operation for interconnected cultural heritage environments. Key takeaways include:

1. **Adaptive Learning Efficiency:** Federated reinforcement learning achieved convergence with minimal communication overhead, enabling scalable deployment.
2. **Cognitive Balance:** The MAPE-K loop successfully maintained equilibrium between conflicting operational goals, confirming its viability as a self-management mechanism.
3. **Sustainability Enhancement:** The system improved energy efficiency by over 150% while preserving artifact stability and visitor comfort.




Figure 6.6: Federated learning reward progression across Museums A–D using PPO + FedProx integrated with NetLogo simulation.

- 4. Resilience and Scalability:** The distributed design ensured robust operation even in the presence of communication delays or environmental variability.

These results provide compelling evidence that CFDT can serve as a foundational model for intelligent, data-driven cultural heritage management systems capable of real-time self-optimization and inter-institutional collaboration. We will publicly release the simulation code, environment wrappers, and trained models under an open license on GitHub upon acceptance. The simulation seeds, parameter files, and scripts to reproduce plots will be included for reproducibility.

6.6 Conclusion and Future Work

This work introduced a Cognitive Federated Digital Twin (CFDT) framework that combines digital twin modeling, federated reinforcement learning, and MAPE-K cognitive control to enable collaborative, privacy-aware optimization for cultural heritage institutions. Through a hybrid NetLogo + Python simulation, we demonstrated that CFDT agents can learn adaptive control policies that balance preservation, comfort, safety, and energy efficiency while sharing knowledge via lightweight federated updates.

Key takeaways are: (i) federated reinforcement learning with FedProx reduces client drift under non-IID conditions; (ii) MAPE-K meta-control improves learning stability and enables dynamic




Figure 6.7: Floor plans for the four simulated museums used in the NetLogo environment. Each museum contains three rooms and a central exhibition space: Museum A (red), Museum B (blue), Museum C (green), Museum D (yellow).

adaptation to environmental drift; and (iii) a hybrid simulation pipeline supports fine-grained agent interactions with a centralized cognitive coordinator without exchanging raw sensor data.

For future work we will: (1) strengthen privacy and robustness by integrating differential privacy and secure aggregation mechanisms; (2) validate CFDT on real museum testbeds with IoT sensors and edge devices; (3) extend the Knowledge component to a semantic knowledge graph for ontology-guided policy transfer across heterogeneous sites; (4) explore model compression and quantization to minimize communication and energy cost (green-AI strategies); and (5) design curator-facing dashboards for human-in-the-loop oversight and safe policy intervention. These directions will move CFDT from a controlled simulation to practical, ethically governed, and energy-efficient deployments for cultural heritage preservation.

Chapter 7

Discussion and Synthesis

This chapter provides a comprehensive discussion and theoretical synthesis of the research conducted throughout this dissertation, which explored the design, implementation, and validation of *Federated and Cognitive Digital Twins (CFDTs)* for cultural heritage. The chapter aims to position the proposed frameworks within broader academic and technological discourses, emphasizing their implications for digital transformation, heritage sustainability, and intelligent systems design. Integrates the findings of the systematic literature review and the empirical results from three successive frameworks Intelligent, Real-Time, and Cognitive Federated Digital Twins demonstrating how each stage contributes to the maturity and scalability of cultural heritage digital ecosystems.

The discussion unfolds in five core analytical dimensions. First, it synthesizes insights across the research stages, from conceptual foundations to cognitive intelligence. Second, it articulates the theoretical implications of the federated paradigm. Third, it explores the technical and operational contributions of the frameworks. Fourth, it evaluates the ethical and cultural implications of adopting intelligent digital twins in heritage contexts. Finally, it addresses interoperability, governance, and policy challenges associated with the adoption of federated digital twin systems.

7.1 Synthesis Across Research Stages

7.1.1 From Systematic Review to Conceptualization

The research began with a systematic literature review of 108 studies on Digital Twins (DTs) in cultural heritage, forming the empirical and conceptual basis of this dissertation. The tri-dimensional analytical framework encompassing user-centric applications, enabling technologies, and maturity levels revealed a fragmented and uneven landscape of digital heritage innovation [18, 30]. While immersive visualization, 3D reconstruction, and documentation dominated the literature, very few studies incorporated real-time data integration, adaptive intelligence, or cross-institutional collaboration. Most DT applications remained static, focusing on visualization rather than predictive or prescriptive modeling [19, 228].

The review identified critical gaps: (1) a lack of interoperability and standardized ontologies, (2) limited incorporation of AI and real-time analytics, (3) the absence of federated or privacy-preserving mechanisms for institutional collaboration, and (4) minimal emphasis on ethical and governance dimensions. These gaps justified the need for a federated, intelligent, and ethically governed DT framework capable of linking heritage institutions through secure and scalable data

sharing mechanisms [139, 280]. Thus, the systematic review not only mapped the state of the art but also directed the subsequent research trajectory toward federated and cognitive digital twins.

7.1.2 Intelligent Federated Digital Twin Framework

Building on the SLR findings, the AI-driven Federated Digital Twin (AIFDT) framework introduced an **intelligent and simulation-validated** framework capable of supporting multi-museum collaboration, distributed decision-making, and privacy-preserving emergency coordination. Unlike the systematic analysis chapter 3, this stage Chapter 4 grounds the framework in a concrete architecture and a detailed agent-based simulation, producing empirical evidence of how federated intelligence can improve operational awareness and response efficiency across interconnected cultural heritage sites.

This stage introduces the Intelligent Federated Digital Twin (IFDT) as a **multi-layered architecture** composed of:

1. a local perception and sensing layer for capturing visitor movement, environmental conditions, and emergency events;
2. a local digital twin model responsible for internal reasoning, hazard assessment, and local prediction;
3. a federated learning layer enabling cross-museum knowledge aggregation using privacy-preserving learning techniques; and
4. an AI-driven decision layer that supports adaptive evacuation routes and operational responses.

Each museum operates as an autonomous Digital Twin instance, retaining full control and ownership of its internal data. Through the federated learning mechanism (FedAvg), museums contribute model updates rather than raw sensor or visitor data, ensuring compliance with organizational privacy requirements. This design is particularly relevant for cultural heritage institutions that frequently face regulatory, ethical, and logistical barriers to sharing sensitive information. This study also presents the first **simulation-based validation** of the federated architecture, implemented using an agent-based evacuation scenario in two heterogeneous museums. The simulation demonstrates that federated intelligence can substantially enhance system-wide performance under emergency conditions such as fires or structural hazards. Quantitative results show improvements in:

- **evacuation time**, with faster overall clearance of affected buildings;
- **congestion reduction**, especially at narrow exits and high-density corridors;
- **cross-museum coordination**, enabled by shared predictive insights; and
- **model accuracy**, with federated learning outperforming isolated local models.

Scenario-driven experiments conducted in Chapter 4 demonstrate how the IFDT architecture performs under varying hazard intensities, visitor distributions, and federated synchronization intervals. These experiments highlight the architecture's scalability and robustness, revealing that

collaborative learning converges more rapidly and produces more stable evacuation strategies than non-federated approaches. The results confirm that federated intelligence can transform localized situational awareness into global, coordinated decision support without compromising privacy.

In addition to showcasing the strengths of the Intelligent FDT framework, Chapter 4 also identifies limitations of non-real-time federated coordination, including communication overhead, potential misalignment across local objective functions, and discrepancies between simulated and real-world conditions. These limitations motivate the enhancements introduced in Chapter 5, where the framework is extended toward more deterministic, real-time collaboration mechanisms.

Overall, this research stage demonstrates that the Intelligent Federated Digital Twin framework is not merely conceptual but a **practical, scalable, and privacy-preserving architecture**. It provides the first quantitative evidence that federated intelligence can significantly enhance the resilience and operational efficiency of cultural heritage institutions, forming a critical bridge between conceptual design and future real-world deployment.

7.1.3 Real-Time Federated Digital Twin Framework

The Real-Time Federated Digital Twin Framework (RT-FDTF) represented the next evolutionary step, addressing temporal synchronization and latency critical factors for emergency response and safety management in heritage contexts. Although the chapter 4 introduced autonomy, it lacked guaranties of real-time determinism. RT-FDTF resolved this through time-triggered scheduling, bounded latency mechanisms, and deterministic synchronization [207].

This framework enabled multi-museum coordination during emergencies such as fire or crowd evacuation, ensuring that decisions were not only intelligent but temporally consistent. Through agent-based modeling and real-time data simulation, the framework achieved an average evacuation success rate exceeding 89% with minimal coordination delays.

In theoretical terms, the RT-FDTF bridged the gap between cyber-physical synchronization and cognitive reasoning. It demonstrated how real-time control theory could coexist with federated learning, providing a blueprint for safety-critical digital heritage systems. Operationally, it validated the feasibility of federated intelligence under deterministic constraints a crucial milestone before advancing toward cognition.

7.1.4 Cognitive Federated Digital Twin Framework

The Cognitive Federated Digital Twin Framework (CFDT) represented the culmination of this research trajectory, integrating semantic reasoning, reinforcement learning, and ethical governance into a unified cognitive architecture. Unlike previous frameworks, the CFDT embodied not only intelligence but understanding an ability to contextualize, reason, and learn beyond immediate operational tasks.

Drawing on cognitive science and knowledge engineering, the CFDT incorporated knowledge graphs and ontologies such as CIDOC CRM to enable semantic interoperability among federated heritage systems [50, 121]. Reinforcement learning agents facilitated continuous policy optimization, while the ethical layer embedded fairness and explainability within decision-making.

The CFDT thus redefined digital twins as autonomous yet ethically aligned cognitive agents within the digital heritage ecosystem. The framework extended beyond real-time optimization

toward collective intelligence where museums collaborate semantically and ethically, maintaining data sovereignty while sharing contextual knowledge.

In synthesis, the CFDT represents the theoretical apex of the dissertation: a holistic integration of perception, intelligence, real-time determinism, and cognition. It fulfills the research vision articulated in Chapter 1 to develop federated, intelligent, and sustainable digital twin ecosystems for cultural heritage preservation.

7.2 Theoretical Implications

The theoretical significance of this research lies in unifying three traditionally separate paradigms: digital twin theory, federated learning, and cognitive intelligence. The synthesis of these paradigms produces a new ontological model for intelligent cultural heritage systems.

First, the research advances **Digital Twin Theory** by moving beyond industrial replication toward socio-technical ecosystems. The dissertation conceptualizes DTs as living cognitive entities that not only mirror but interpret and influence their physical counterparts.

Second, it extends **Federated Learning Theory** from data-centric optimization to organizational intelligence. Here, federated models are not mere statistical aggregators but mediators of institutional collaboration and ethical autonomy.

Third, the dissertation contributes to **Cognitive Systems Theory** by embedding ethical reasoning, semantic alignment, and self-adaptation within the digital twin lifecycle. This integration establishes the foundation for “Cognitive Federated Twins,” a new class of intelligent systems capable of sustainable cultural stewardship. These theoretical contributions collectively redefine digital heritage informatics as a discipline grounded in distributed cognition, ethical AI, and systemic interoperability.

7.3 Ethical and Cultural Dimensions

A significant dimension of this research is the ethical and cultural responsibility embedded in the deployment of intelligent systems within heritage contexts. Unlike industrial or urban applications, cultural heritage carries intrinsic values related to identity, memory, and authenticity. Therefore, any technological intervention must adhere to ethical and human-centered principles that respect cultural diversity, intellectual property, and community participation [29, 69, 251].

The introduction of federated and cognitive digital twins into cultural heritage institutions raises critical questions about data governance, transparency, and inclusivity. The federated approach adopted in this dissertation mitigates ethical risks by maintaining data sovereignty ensuring that each museum or cultural institution retains control over its data while participating in collective learning. This aligns with the principles of decentralized autonomy and ethical AI governance articulated by UNESCO (2022).

The cognitive layer further enhances ethical compliance through explainable and interpretable AI. Reinforcement learning agents were designed with human-in-the-loop validation to prevent algorithmic bias and to align decision outcomes with curatorial objectives. Moreover, semantic reasoning facilitates transparent provenance tracking and accountability across institutions. By embedding

ethical design at the core of technical architecture, this research demonstrates that advanced AI systems can serve as custodians rather than disruptors of cultural authenticity.

Culturally, the proposed frameworks promote resilience, inclusivity, and sustainability. The integration of XR technologies not only preserves tangible and intangible heritage but also democratizes access to cultural knowledge. Through immersive visualization and remote participation, marginalized communities and global audiences can experience cultural assets in equitable ways. This aligns with the UN Sustainable Development Goals (SDG 4, 9, and 11), particularly in promoting inclusive education, innovation, and sustainable cities [252].

The dissertation thus reinforces the notion that technology must act as an enabler of cultural continuity. By embedding ethics, transparency, and inclusivity in its architecture, the Federated and Cognitive Digital Twin paradigm ensures that digital transformation strengthens cultural value rather than commodifying it.

7.3.1 Data Quality, Provenance, and FAIR/CARE Compliance

The credibility and long-term sustainability of Cultural Heritage Digital Twins strongly depend on the quality, authenticity, provenance, and trustworthiness of their associated information resources. Since Digital Twins increasingly integrate heterogeneous data originating from historical archives, sensor streams, semantic repositories, community contributions, and immersive environments, rigorous governance mechanisms are required to preserve data reliability and contextual integrity. To address these challenges, the proposed framework aligns with both FAIR (Findable, Accessible, Interoperable, and Reusable) and CARE (Collective Benefit, Authority to Control, Responsibility, and Ethics) principles for heritage data management [37, 242]. From the FAIR perspective, all digital assets and related knowledge resources should include: persistent identifiers, machine-readable metadata, semantic interoperability standards, and provenance-aware documentation mechanisms. These capabilities improve resource discoverability, interoperability between federated systems, and long-term reusability across Cultural Heritage institutions. Simultaneously, CARE principles ensure that the management and dissemination of heritage knowledge remain ethically responsible and culturally sensitive. In particular, Digital Twins should respect community ownership, cultural authority, indigenous rights, and access-control requirements associated with sensitive or sacred heritage information [186].

7.4 Interoperability and Policy Implications

Interoperability represents both a technical and institutional challenge in the digital heritage domain. The systematic review identified fragmentation as one of the key obstacles to the adoption of advanced digital twin systems 3. Different heritage organizations employ heterogeneous data formats, metadata standards, and digital platforms, leading to isolated digital silos.

The federated architecture developed in this dissertation addresses interoperability through a multi-layered approach. At the technical level, semantic interoperability was achieved using standardized ontologies such as CIDOC CRM, ensuring that heritage data could be semantically linked between institutions, as clearly illustrated in the chapter 4. At the organizational level, federated learning enabled collaborative model training without compromising privacy or data ownership, creating a shared intelligence layer across distributed museums. From a policy perspective, the

adoption of federated frameworks introduces a new governance model for cultural heritage data. Rather than centralizing digital heritage repositories, which can expose institutions to data security and ownership risks, federated networks promote distributed governance. Each node (museum, archive, or cultural site) acts as an autonomous participant in a broader ecosystem of collaborative intelligence.

This decentralized paradigm aligns with emerging digital sovereignty frameworks in Europe and worldwide [29, 250]. Governments and heritage agencies can leverage federated architectures to establish secure, ethical, and interoperable infrastructures for cultural data exchange. Additionally, the maturity assessment model proposed in this dissertation can serve as a policy tool for evaluating the digital readiness of heritage institutions, helping policymakers allocate resources more strategically.

In essence, the interoperability and policy contributions of this dissertation extend beyond technical feasibility. They provide a roadmap for structuring future cultural heritage ecosystems around the principles of sovereignty, ethics, and collaboration.

Chapter 8

Conclusion and Future Work

This study has advanced the state of the art in Digital Twin (DT) technologies for cultural heritage by introducing a progressive evolution from static digital replicas to Federated Digital Twins (FDTs) and ultimately to Cognitive Federated Digital Twins (CFDTs). The research addressed critical challenges of interoperability, privacy, and adaptive intelligence through a multi-stage approach that integrated systematic analysis, architectural design, simulation-based validation, and cognitive modeling.

The first major contribution is the systematic review of Digital Twin applications in cultural heritage. This review provided a tri-dimensional analytical framework encompassing user-centric applications, enabling technologies, and maturity assessment. It revealed that most existing implementations remain at an early maturity level, characterized by static models and limited real-time interaction, thereby justifying the need for federated and cognitive extensions.

The second contribution is the development of an AI-driven Federated Digital Twin (FDT) architecture for decentralized cultural heritage digitization. This framework demonstrated how federated learning can enable collaborative intelligence across multiple institutions without compromising data sovereignty. By integrating semantic metadata standards and privacy-preserving protocols, the proposed architecture aligns with ethical and legal requirements such as GDPR and UNESCO's AI ethics guidelines.

The third contribution is the Real-Time Federated Digital Twin Integration Framework (FRT-DTIF), which introduced deterministic scheduling and Worst-Case Execution Time (WCET) profiling into federated architectures. Simulation-based validation confirmed that the framework can maintain bounded synchronization latency and execute decision cycles within a 500 ms deadline, enabling reliable coordination in safety-critical emergency scenarios.

The final and most advanced contribution is the Cognitive Federated Digital Twin (CFDT) model. By embedding reinforcement learning and the MAPE-K autonomic control loop, the CFDT framework introduced self-learning capabilities that enable continuous adaptation to dynamic environmental conditions. Experimental results demonstrated significant improvements in energy efficiency, artifact preservation, and visitor comfort, validating the hypothesis that cognitive intelligence can transform digital twins from reactive systems into proactive, knowledge-driven entities.

8.1 Scientific and Practical Implications

The findings of this research have profound implications for both theory and practice. From a theoretical perspective, the integration of federated learning and cognitive adaptation within digital twin architectures advances the conceptual boundaries of cyber-physical systems. It demonstrates that distributed intelligence and real-time coordination can coexist under strict timing and privacy constraints, paving the way for future research in edge-based AI and autonomous decision-making.

From a practical standpoint, the proposed frameworks offer a scalable and ethically grounded solution for cultural heritage management. By enabling decentralized collaboration, real-time responsiveness, and adaptive control, these systems can enhance operational resilience, improve visitor safety, and support sustainable preservation strategies. Furthermore, the alignment with international standards ensures that the proposed solutions are not only technologically robust but also socially responsible.

This thesis additionally highlighted the importance of stakeholder-oriented digital twin design principles, multidisciplinary user requirements, and FAIR/CARE-compliant heritage knowledge governance. The proposed federated and intelligent framework therefore contributes not only to the technological advancement of digital twins, but also to their evolution into sustainable and trustworthy knowledge ecosystems capable of preserving both tangible and intangible Cultural Heritage.

8.2 Limitations

Despite its contributions, this research is not without limitations. The validation relied primarily on simulation environments, which, although ethically responsible, cannot fully replicate real-world complexities such as network jitter, heterogeneous hardware, and unpredictable human behavior during emergencies. Additionally, the scalability of federated learning under extreme heterogeneity and the integration of immersive technologies such as Extended Reality (XR) require empirical testing. Finally, while ethical compliance was considered at the architectural level, operationalizing these principles in live deployments demands further research.

8.3 Future Work

Building on the insights and limitations identified, future research will pursue five strategic directions to advance the proposed frameworks toward real-world applicability and broader impact.

8.3.1 Technical Enhancements and AI Integration

Future work will incorporate advanced federated optimization algorithms such as FedProx, SCAFFOLD, and adaptive asynchronous strategies to improve convergence under non-IID conditions. Additionally, the integration of uncertainty modeling and probabilistic decision-making will enable the system to handle sensor noise and dynamic environmental fluctuations more effectively. Research will also explore adaptive AI models for predictive maintenance and personalized evacuation planning, leveraging historical and real-time data streams to enhance decision accuracy and responsiveness.

8.3.2 Immersive Interfaces and Spatial Realism

A second priority is the extension of the CFDT framework with immersive technologies that enhance spatial realism and stakeholder engagement. Building Information Modeling (BIM), LiDAR-based 3D models, and Extended Reality (XR) interfaces will be integrated to support interactive training simulations and participatory planning. These enhancements will allow curators, emergency responders, and visitors to experience dynamic scenarios in virtual environments, thereby improving preparedness and fostering collaborative decision-making.

8.3.3 Real-World Deployment and Validation

The third research direction focuses on transitioning from simulation-based validation to empirical testing in operational environments. The current prototype corresponds to Technology Readiness Level (TRL) 4; future work aims to achieve TRL 6 through pilot implementations in museums and archaeological sites. This will involve hardware-in-the-loop (HIL) testing using edge devices such as Raspberry Pi and NVIDIA Jetson, integrated with IoT sensors for real-time data acquisition. Empirical benchmarking under realistic network conditions including latency variability, packet loss, and bandwidth constraints will provide critical insights into system robustness and scalability.

8.3.4 Ethical Governance and Open Science

Ensuring ethical compliance and transparency remains a cornerstone of future research. Privacy-preserving techniques such as differential privacy and secure aggregation will be implemented within federated learning workflows to strengthen data protection. Furthermore, stakeholder-informed governance models will be developed to guide data sharing, model training, and decision-making processes across institutions. To promote reproducibility and community engagement, simulation models, architectural blueprints, and selected software modules will be released under open-source licenses, aligning with FAIR data principles and the broader movement toward open science.

8.3.5 Scalability and Sustainability

The final research direction addresses the twin challenges of scalability and environmental sustainability. Computational efficiency will be optimized through model compression techniques and green AI strategies aimed at reducing energy consumption. Semantic interoperability will be extended using linked data and knowledge graphs to facilitate cross-institutional policy transfer and collaborative analytics. These efforts will enable the CFDT framework to scale beyond individual museums to large cultural ecosystems and smart city infrastructures, positioning digital twins as key enablers of sustainable urban heritage management.

Final Reflection

In conclusion, this dissertation contributes a foundational approach to the design and implementation of intelligent digital ecosystems for cultural heritage. By bridging the gap between theoretical innovation and practical application, it lays the groundwork for a future where cultural heritage

preservation evolves from static documentation toward dynamic, collaborative, and sustainable digital infrastructures. This trajectory not only safeguards humanity's collective memory but also redefines how we engage with heritage in an increasingly interconnected and uncertain world.

Bibliography

- [1] M. Abdelalim. Heritage preservation using laser scanning: Architectural digital twins using al-mu'izz street as a case study. *Buildings*, 15(9):1480, 2025.
- [2] M. A. Adibhesami, H. Karimi, B. Sepehri, F. Pajani, S. Makaremi, S. Ghahremani, and A. Hassanzadeh. A data-driven multi-scale digital twin framework for optimizing energy efficiency in public pedestrian infrastructure. In *Digital Twin Computing for Urban Intelligence*, pages 147–166. Springer, 2024.
- [3] C. Alcaraz and J. Lopez. Real-time requirements in smart systems: A review of research and challenges. *Future Generation Computer Systems*, 117:288–302, 2021.
- [4] C. Alcaraz and J. Lopez. Digital twin: A comprehensive survey of security threats. *IEEE Communications Surveys & Tutorials*, 24(3):1475–1503, 2022.
- [5] K. Alfalqi and M. Bellaiche. Emergency events detection based on integration of federated learning and active learning. *International Journal of Information Technology*, 15(6):2863–2876, 2023.
- [6] M. Almatared, H. Liu, O. Abudayyeh, O. Hakim, and M. Sulaiman. Digital-twin-based fire safety management framework for smart buildings. *Buildings*, 14(1):4, 2023.
- [7] A. Amelio and G. P. Zarri. Cultural heritage digital twin: Modeling and representing the visual narrative in leonardo da vinci's mona lisa. *Neural Computing and Applications*, 36(20):11859–11876, 2024.
- [8] S. B. Amor, N. Elloumi, A. Eltaief, B. Louhichi, N. H. Alrasheedi, and A. Seibi. Digital twin implementation in additive manufacturing: A comprehensive review. *Processes*, 12(6):1062, 2024.
- [9] L.-M. Angheluță, A. Ignuța Acimov, C. Gora, A. I. Chiricuță, A. I. Popovici, and V. Obradovici. Documenting romania's wooden churches: Integrating modern digital platforms with vernacular conservation. *Heritage*, 8(3):103, 2025.
- [10] M. S. Anwar, J. Yang, J. Frnda, A. Choi, N. Baghaei, and M. Ali. Metaverse and xr for cultural heritage education: applications, standards, architecture, and technological insights for enhanced immersive experience. *Virtual Reality*, 29(2):51, 2025.
- [11] F. I. Apollonio, F. Fantini, S. Garagnani, and M. Gaiani. A photogrammetry-based workflow for the accurate 3d construction and visualization of museums assets. *Remote Sensing*, 13(3):486, 2021.

- [12] A. Bajwa. Ai-based emergency response systems: A systematic literature review on smart infrastructure safety. *Available at SSRN 5171521*, 2025.
- [13] F. Banfi. The evolution of interactivity, immersion and interoperability in hbim: Digital model uses, vr and ar for built cultural heritage. *ISPRS International Journal of geo-information*, 10(10):685, 2021.
- [14] J. Banks, J. Carson, B. Nelson, and D. Nicol. *Discrete-Event System Simulation*. Pearson, 5th edition, 2010.
- [15] P. M. Barone. Between past and future: Digital technologies and the revolution in cultural heritage preservation. *DigItalia*, 19(2):81–99, 2024.
- [16] A. Bašić and Others. Promoting cultural heritage through a micro-business by means of a digital twin. In *Proceedings of the 2024 International Conference on Cultural Heritage and Digital Innovation*, pages 1–8. ResearchGate, 2024.
- [17] G. Beale and et al. Inclusive digital heritage: Opportunities and challenges in participatory practice. *Journal of Cultural Heritage Management and Sustainable Development*, 2022.
- [18] M. K. Bekele and E. Champion. A survey of augmented, virtual, and mixed reality for cultural heritage. *Journal on Computing and Cultural Heritage*, 11(2):1–36, 2018.
- [19] M. K. Bekele and E. Champion. A review of immersive virtual heritage. *Frontiers in Virtual Reality*, 2:16, 2021.
- [20] M. R. A. Berkani, A. Chouchane, Y. Himeur, A. Ouamane, S. Miniaoui, S. Atalla, W. Mansoor, and H. Al-Ahmad. Advances in federated learning: Applications and challenges in smart building environments and beyond. *Computers*, 14(4):124, 2025.
- [21] C. Bertolin. Preservation of cultural heritage and resources threatened by climate change, 2019.
- [22] L. Bezzi, A. Bezzi, R. Gietl, C. Moraes, G. Naponiello, S. Airò, A. Putzer, and E. Silvestri. Digital twins of archaeological finds: Open source technologies applied to 3d scanning. *Archeologia e Calcolatori*, 35(2):203–214, 2024.
- [23] C. Binding. Implementing archaeological time periods using cidoc crm and skos. In *Extended Semantic Web Conference*, pages 273–287. Springer, 2010.
- [24] C. M. Bolognesi and F. Fiorillo. Virtual representations of cultural heritage: Sharable and implementable case study to be enjoyed and maintained by the community. *Buildings*, 13(2):410, 2023.
- [25] G. Bonari, A. D’Onofrio, S. Cernaianu, et al. Digital twins for urban contaminant diffusion: A computational approach. *IEEE Access*, 12:20056–20070, 2024.
- [26] S. Bond and D. Worthing. *Managing built heritage: The role of cultural values and significance*. John Wiley & Sons, 2016.

- [27] L. Bonorden, M. Frerichs, M. Riebisch, S. von Riegen, F. Hartke, R. Herzog, L. Hotz, D. Jürgensen, M. Kiele-Dunsche, S. Schottler, et al. Decision-making about federated digital twins—how to distribute information storage and computing. In *Modellierung 2022*, pages 49–64. Gesellschaft für Informatik eV, 2022.
- [28] E. Boyko, T. Volkova, M. Kucher, V. Minenkova, D. Komarov, and N. Pikalova. Creating 3d models as digital twins of objects of the cultural and historical heritage of krasnodar krai. *Revista Gestão & Tecnologia*, 24(Special):222–232, 2024.
- [29] P. Brunet, L. De Luca, E. Hyvönen, and G. Sonkoly. Report on a european collaborative cloud for cultural heritage: ex-ante impact assessment, 2022.
- [30] F. Bruno, S. D’Auria, M. Gallo, R. Santoro, and G. Saponaro. Digital twins for cultural heritage: Concepts and applications. *Digital Applications in Archaeology and Cultural Heritage*, 25:e00222, 2022.
- [31] S. Bruno, A. Sciotti, A. Pierucci, R. Rubino, T. Di Noia, and F. Fatiguso. Verbum—virtual enhanced reality for building modelling (virtual technical tour in digital twins for building conservation). *Journal of Information Technology in Construction*, 27, 2022.
- [32] A. Bucciero, A. Chirivì, R. Colella, M. Emara, M. Greco, D. M. Palamà, A. Pandurino, F. Taurino, and D. Zecca. H2iosc project: The italian federated cluster for iot-based monitoring and digital twinning of cultural heritage. In *2024 9th International Conference on Smart and Sustainable Technologies (SpliTech)*, pages 1–6. IEEE, 2024.
- [33] G. Buttazzo. *Hard Real-Time Computing Systems: Predictable Scheduling Algorithms and Applications*. Springer, 2011.
- [34] D. Calvaresi, M. Marinoni, A. Sturm, M. Schumacher, and G. Buttazzo. The challenge of real-time multi-agent systems for enabling iot and cps. In *Proceedings of the international conference on web intelligence*, pages 356–364, 2017.
- [35] A. Calzado-Martínez, Á.-L. García-Fernández, L. M. Ortega-Alvarado, and F.-R. Feito-Higueruela. Integrated information system for 3d interactive reconstruction of an archaeological site. *ACM Journal on Computing and Cultural Heritage*, 16(3):1–23, 2023.
- [36] C. Capurro, G. Plets, and J. Verheul. Digital heritage infrastructures as cultural policy instruments: Europeana and the enactment of european citizenship. *International Journal of Cultural Policy*, 30(3):304–324, 2024.
- [37] S. R. Carroll, I. Garba, O. L. Figueroa-Rodríguez, J. Holbrook, R. Lovett, S. Materechera, M. Parsons, K. Raseroka, D. Rodriguez-Lonebear, R. Rowe, et al. The care principles for indigenous data governance. 2020.
- [38] A. Castiglione, F. Colace, V. Moscato, and F. Palmieri. Chis: A big data infrastructure to manage digital cultural items. *Future Generation Computer Systems*, 86:1134–1145, 2018.
- [39] L. Cecere, F. Colace, G. D’Aniello, M. Gaeta, M. D. Santo, and S. D. Paola. Predictive maintenance of an archaeological park: An iot and digital twin based approach. In R. S.

- Chaves and M. F. Santos, editors, *Artificial Intelligence in Human-Computer Interaction*, pages 323–341. Springer, 2024.
- [40] L. Cecere, F. Colace, A. Lorusso, B. Messina, A. Tucker, and D. Santaniello. Iot and digital twin: A new perspective for cultural heritage predictive maintenance. *Procedia Structural Integrity*, 64:2181–2188, 2024.
- [41] P. Centorrino, A. Corbetta, E. Cristiani, and E. Onofri. Managing crowded museums: Visitors flow measurement, analysis, modeling, and optimization. *Journal of Computational Science*, 53:101357, 2021.
- [42] E. Champion. *Critical Gaming and Virtual Heritage*. Routledge, 2021.
- [43] J. Chen. Optimizing lidar point clouds for mobile digital twin restoration of cultural heritage with trustworthy validation. *International Journal of Interactive Mobile Technologies*, 19(9), 2025.
- [44] L. Chen, X. Xie, Q. Lu, A. K. Parlikad, M. Pitt, and J. Yang. Gemini principles-based digital twin maturity model for asset management. *Sustainability*, 13(15):8224, 2021.
- [45] W. Chen, M. Li, and X. Zhao. Real-time digital twin scheduling in smart manufacturing. *Robotics and Computer-Integrated Manufacturing*, 84:102542, 2023.
- [46] J. Cheng, H. Zhang, F. Tao, and C.-F. Juang. Dt-ii: Digital twin enhanced industrial internet reference framework towards smart manufacturing. *Robotics and Computer-Integrated Manufacturing*, 62:101881, 2020.
- [47] S. Chernbumroong, P. Ariya, S. Yolthasart, N. Wongwan, K. Intawong, and K. Puritat. Comparing the impact of non-gamified and gamified virtual reality in digital twin virtual museum environments: A case study of wieng yong house museum, thailand. *Heritage*, 7(4):1870–1892, 2024.
- [48] S. R. Chhetri, S. Faezi, A. Canedo, and M. A. A. Faruque. Quilt: Quality inference from living digital twins in iot-enabled manufacturing systems. In *Proceedings of the international conference on internet of things design and implementation*, pages 237–248, 2019.
- [49] E. Ch’ng. The first original copy and the role of blockchain in the reproduction of cultural heritage. *Presence*, 27(1):151–162, 2019.
- [50] CIDOC CRM. Cidoc conceptual reference model. <https://www.cidoc-crm.org/>, 2007.
- [51] F. Cinquepalmi and F. Cumo. Using digital twin models (dtm) for managing, protecting, and restoring historical buildings. *Conservation Science in Cultural Heritage*, 22:425–445, 2022.
- [52] F. Cinquepalmi and F. Cumo. Using digital twin models (dtm) for managing, protecting and restoring historical buildings. *Conservation Science in Cultural Heritage*, 22:425–445, 2022.
- [53] H. Cooper, L. V. Hedges, and J. C. Valentine. *The handbook of research synthesis and meta-analysis*. Russell Sage Foundation, 2019.

- [54] M. Ćosović and M. Maksimović. Application of the digital twin concept in cultural heritage. In *Proceedings of VIPERC2022: 1st International virtual conference on visual pattern extraction and recognition for cultural heritage understanding*, volume 12, 2022.
- [55] P. A. Cruz Franco, A. Rueda Márquez de la Plata, and E. Gómez Bernal. Protocols for the graphic and constructive diffusion of digital twins of the architectural heritage that guarantee universal accessibility through ar and vr. *Applied Sciences*, 12(17):8785, 2022.
- [56] G. A. Dagnaw, R. Capuano, and H. Muccini. Replication package - digital twins for cultural heritage: a systematic analysis of the state of the art. <https://osf.io/np94b/files/osfstorage>, 2025.
- [57] G. A. Dagnaw and H. Muccini. Ai-powered federated digital twins for enhanced cultural heritage preservation and predictive analytics. *ERCIM News*, (141), April 2025.
- [58] M. Dalibor, N. Jansen, B. Rumpe, D. Schmalzing, L. Wachtmeister, M. Wimmer, and A. Wortmann. A cross-domain systematic mapping study on software engineering for digital twins. *Journal of Systems and Software*, 193:111361, 2022.
- [59] R. Damiano, V. Lombardo, A. Lieto, and D. Borra. Mind the red thread! 3d metaphors for cultural heritage visualization. In *2015 Digital Heritage*, volume 2, pages 355–358. IEEE, 2015.
- [60] X. Dang, W. Liu, Q. Hong, Y. Wang, and X. Chen. Digital twin applications on cultural world heritage sites in china: A state-of-the-art overview. *Journal of Cultural Heritage*, 64:228–243, 2023.
- [61] D. S. de Almeida, I. Boavida-Portugal, et al. Digital twins in tourism: a systematic literature review. *arXiv preprint arXiv:2502.00002*, 2025.
- [62] R. G. De Paoli, E. Di Miceli, and F. Giuliani. Disasters and cultural heritage: planning for prevention, emergency management and risk reduction. In *IOP Conference Series: Materials Science and Engineering*, volume 949, page 012084. IOP Publishing, 2020.
- [63] L. A. Dhahi and S. C. Studio. Louvre abu dhahi launches quantum dome project, amplifying visitor experience through virtual reality technology. 2025.
- [64] R. Di Giulio. Digital transition strategies for architectural heritage. In *ETHICS: Endorse Technologies for Heritage Innovation: Cross-Disciplinary Strategies*, pages 133–149. Springer, 2024.
- [65] U. S. Diego. Digital reconstruction of ancient temple on display in uc san diego library’s art of science exhibit. 2023. Accessed: 2025-08-30.
- [66] M. S. Dihan, A. I. Akash, Z. Tasneem, P. Das, S. K. Das, M. R. Islam, M. M. Islam, F. R. Badal, M. F. Ali, M. H. Ahamed, et al. Digital twin: Data exploration, architecture, implementation and future. *Heliyon*, 10(5), 2024.
- [67] Y. Ding, Y. Zhang, and X. Huang. Intelligent emergency digital twin system for monitoring building fire evacuation. *Journal of Building Engineering*, 77:107416, 2023.

- [68] M. Doerr. The cidoc conceptual reference model: An ontological approach to semantic interoperability of metadata. *AI Magazine*, 24(3), 75–92, 2003.
- [69] M. Donnelly. Cultural data governance and ethics. *Journal of Information Ethics*, 29(2):12–25, 2020.
- [70] O. T. Eleftheriou and C.-N. Anagnostopoulos. Digital twins: A brief overview of applications, challenges and enabling technologies in the last decade. *Digital Twin*, 2:2, 2022.
- [71] M. Emara, M. Castellini, C. Riminesi, M. A. Jaziri, A. Pandurino, A. Conti, L. Fiorini, A. Meucci, R. Colella, F. Taurino, et al. Iot-driven digital twin based on vr for smart monitoring of cultural heritage sites: La grotta degli animali in florence case study. In *International Conference on Extended Reality*, pages 496–515. Springer, 2025.
- [72] M. Falcone, A. Origlia, M. Campi, and S. Di Martino. From architectural survey to continuous monitoring: Graph-based data management for cultural heritage conservation with digital twins. *The International Archives of the Photogrammetry, Remote Sensing and Spatial Information Sciences*, 43:47–53, 2021.
- [73] E. Farella, L. Morelli, E. Grilli, S. Rigon, and F. Remondino. Handling critical aspects in massive photogrammetric digitization of museum assets. *The International Archives of the Photogrammetry, Remote Sensing and Spatial Information Sciences*, 46:215–222, 2022.
- [74] A. Felicetti, A. Himmiche, and M. Somenzi. Artificial intelligence and ontologies for the management of heritage digital twins data. *Journal of Data and Information Science*, 10(1):1, 2024.
- [75] D. Fenoglio, D. Josifovski, A. Gobetti, M. Formo, H. Gjoreski, M. Gjoreski, and M. Langheinrich. Federated learning for privacy-aware cognitive workload estimation. In *Proceedings of the 22nd International Conference on Mobile and Ubiquitous Multimedia (MUM)*, pages 25–36. ACM, 2023.
- [76] L. Fregonese, N. Giordani, A. Adami, G. Bachinsky, L. Taffurelli, O. Rosignoli, and J. Helder. Physical and virtual reconstruction for an integrated archaeological model: 3d print and maquette. *The International Archives of the Photogrammetry, Remote Sensing and Spatial Information Sciences*, 42:481–487, 2019.
- [77] B. Frischer and M. Abdelaziz. Rome reborn: Status report and future of the project. In *Eurographics Association – Digital Heritage Conference Workshop*, page —. Eurographics, 2025.
- [78] A. Fuller, Z. Fan, C. Day, and C. Barlow. Digital twin: enabling technologies, challenges and open research. *IEEE access*, 8:108952–108971, 2020.
- [79] F. Gabellone. Digital twin: a new perspective for cultural heritage management and fruition. *Acta Imeko*, 11(1):7–pp, 2022.
- [80] F. Gabellone. Digital twin: a new perspective for cultural heritage management and fruition. *Acta IMEKO*, 11(1):7–pp, 2022.

- [81] S. Garagnani, R. D’Ugo, A. Lupi, B. Martini, M. Salvucci, M. Susta, M. Tombolato, and R. Barsanti. Visual perception and cognition by the means of interactive digital replicas of museum artifacts: Leonardo da vinci’s drawings as if they were in visitors’ hands. *Heritage*, 6(1):1–25, 2022.
- [82] M. Gerges, P. Demian, and Z. Adamu. Customising evacuation instructions for high-rise residential occupants to expedite fire egress: Results from agent-based simulation. *Fire*, 4(2):21, 2021.
- [83] J. Gil and P. Lambert. Federated data ecosystems for urban digital twins. *Journal of Urban Technology*, 31(2):87–104, 2024.
- [84] L. Gobeawan, E. Lin, A. Tandon, A. Yee, V. Khoo, S. Teo, S. Yi, C. Lim, S. Wong, D. Wise, et al. Modeling trees for virtual singapore: From data acquisition to citygml models. *The International Archives of the Photogrammetry, Remote Sensing and Spatial Information Sciences*, 42:55–62, 2018.
- [85] J. González, K. Figueiredo, A. W. Hammad, V. W. Tam, A. N. Haddad, and C. Illankoon. Heritage bim (hbim) applied in emergency scenarios: A case study of the national museum in brazil. *International Journal of Construction Management*, 25(11):1239–1253, 2025.
- [86] M. Grieves. Digital twin: manufacturing excellence through virtual factory replication. *White paper*, 1(2014):1–7, 2014.
- [87] M. Grieves and J. Vickers. Origins of the digital twin concept. *Florida Institute of Technology*, 8:3–20, 2016.
- [88] M. Grieves and J. Vickers. Digital twin: Mitigating unpredictable, undesirable emergent behavior in complex systems. *Transdisciplinary perspectives on complex systems: New findings and approaches*, pages 85–113, 2017.
- [89] E. Grifoni, L. Bonizzoni, M. Gargano, J. Melada, N. Ludwig, S. Bruni, and I. Mignani. Hyper-dimensional visualization of cultural heritage: A novel multi-analytical approach on 3d pomological models in the collection of the university of milan. *ACM Journal on Computing and Cultural Heritage (JOCCH)*, 15(2):1–15, 2022.
- [90] M. Guarneri, M. F. De Collibus, M. Francucci, and M. Ciaffi. The importance of artworks 3d digitalization at the time of covid epidemy: case studies by the use of a multi-wavelengths technique. In *2020 IEEE 5th International Conference on Image, Vision and Computing (ICIVC)*, pages 113–117. IEEE, 2020.
- [91] Y. Guo, L. Liu, W. Huang, M. Shen, X. Yi, J. Zhang, et al. Extending x-reality technologies to digital twin in cultural heritage risk management: a comparative evaluation from the perspective of situation awareness. *Heritage Science*, 12(1):245, 2024.
- [92] S. Gwynne, E. R. Galea, M. Owen, P. J. Lawrence, and L. Filippidis. A review of the methodologies used in the computer simulation of evacuation from the built environment. *Building and environment*, 34(6):741–749, 1999.

- [93] M. Haibt. End-to-end digital twin creation of the archaeological landscape in uruk-warka (iraq). *International Journal of Digital Earth*, 17(1):2324964, 2024.
- [94] Z. A. Hammadeh, S. Quinton, and R. Ernst. Weakly-hard real-time guarantees for earliest deadline first scheduling of independent tasks. *ACM Transactions on Embedded Computing Systems (TECS)*, 18(6):1–25, 2019.
- [95] X. Hao and Y. Wang. Application of digital twin and meta-universe technology in museum exhibition design. In *Proceedings of the 2024 International Conference on Telecommunication and Electronics (TELEPE)*, pages 262–265. IEEE, 2024.
- [96] E. Hedrick, D. Hatch, and E. Oliver. Digital twins in the ancient and historical world. In *2025 Intermountain Engineering, Technology and Computing (IETC)*, pages 1–6. IEEE, 2025.
- [97] C. Heritage. A digital twin as blueprint for restoring the notre-dame. 2024. Accessed: 2025-08-30.
- [98] S. Hermon, F. Niccolucci, N. Bakirtzis, and S. Gasanova. A heritage digital twin ontology-based description of giovanni baronzio’s ‘crucifixion of christ’ analytical investigation. *Journal of Cultural Heritage*, 66:48–58, 2024.
- [99] R. E. Hollinger, E. John Jr, H. Jacobs, L. Moran-Collins, C. Thome, J. Zastrow, A. Metallo, G. Waibel, and V. Rossi. Tlingit-smithsonian collaborations with 3d digitization of cultural objects. *Museum Anthropology Review*, 7(1-2):201–253, 2013.
- [100] H. Hosamo and S. Mazzetto. Integrating knowledge graphs and digital twins for heritage building conservation. *Buildings*, 15(1):16, 2024.
- [101] H. Hou, J. H. Lai, H. Wu, and T. Wang. Digital twin application in heritage facilities management: systematic literature review and future development directions. *Engineering, Construction and Architectural Management*, 2023.
- [102] D. Hu and J. S. Minner. Uavs and 3d city modeling to aid urban planning and historic preservation: A systematic review. *Remote Sensing*, 15(23):5507, 2023.
- [103] J. Hutson, J. Weber, and A. Russo. Digital twins and cultural heritage preservation: a case study of best practices and reproducibility in chiesa dei ss apostoli e biagio. *Art and Design Review*, 11, 2023.
- [104] C. Ihle, D. Trautwein, M. Schubotz, N. Meuschke, and B. Gipp. Incentive mechanisms in peer-to-peer networks — a systematic literature review. *ACM Comput. Surv.*, 55(14s), July 2023.
- [105] M.-E. Iliuță, M.-A. Moisescu, E. Pop, A.-D. Ionita, S.-I. Caramihai, and T.-C. Mitulescu. Digital twin—a review of the evolution from concept to technology and its analytical perspectives on applications in various fields. *Applied Sciences*, 14(13):5454, 2024.
- [106] D. Isern, A. Moreno, and D. Sánchez. Agents and simulation for emergency situations in cultural heritage buildings. *Expert Systems with Applications*, 41(6):2521–2535, 2014.

- [107] R. Ivanov and V. Velkova. Analyzing visitor behavior to enhance personalized experiences in smart museums: A systematic literature review. *Computers*, 14(5):191, 2025.
- [108] N. Jelenić and N. Ražnatović. Geometric digital twin for conservation project: The case of the st. basil of ostrog church in nikšić (montenegro). *Scientific Culture*, 10(2):81–97, 2024.
- [109] D.-Y. Jeong, M.-S. Baek, T.-B. Lim, Y.-W. Kim, S.-H. Kim, Y.-T. Lee, W.-S. Jung, and I.-B. Lee. Digital twin: Technology evolution stages and implementation layers with technology elements. *Ieee Access*, 10:52609–52620, 2022.
- [110] H. Jia and J. Yan. Construction of heritage digital resource platform based on digital twin technology. *Mathematical Problems in Engineering*, 2022, 2022.
- [111] H. Jiang, S. Qin, J. Fu, J. Zhang, and G. Ding. How to model and implement connections between physical and virtual models for digital twin application. *Journal of Manufacturing Systems*, 58:36–51, 2021.
- [112] C. Joon-cheol. *Worst-case Performance Analysis Techniques for Distributed Real-Time Embedded Systems*. PhD thesis, , 2018.
- [113] P. Jouan and P. Hallot. Digital twin: Research framework to support preventive conservation policies. *ISPRS International Journal of Geo-Information*, 9(4):228, 2020.
- [114] V. Jover and S. Sempere. Creating digital twins to celebrate commemorative events in the metaverse. *Computers*, 14(7):273, 2025.
- [115] P. Kairouz, H. McMahan, and et al. Advances and open problems in federated learning. *Foundations and Trends in Machine Learning*, 14(1–2):1–210, 2021.
- [116] K. Karatzas et al. Digital twins-enabled heritage buildings management through social dynamics. *Journal of Cultural Heritage Management and Sustainable Development*, 14(1):1–21, 2024.
- [117] M. Kaur, S. Singh, A. Sharma, and A. Kumar. Digital twins for smart manufacturing: Architecture and applications. *Procedia CIRP*, 93:234–239, 2020.
- [118] S. Kazanis, G. Kontogianni, R. Chliverou, and A. Georgopoulos. Developing a virtual museum for the ancient wine trade in eastern mediterranean. *The International Archives of the Photogrammetry, Remote Sensing and Spatial Information Sciences*, 42:399–405, 2017.
- [119] Y.-J. Kim, H. Kim, B. Ha, and W.-T. Kim. Advanced fire emergency management based on potential fire risk assessment with informative digital twins. *Automation in Construction*, 167:105722, 2024.
- [120] B. Kitchenham, S. Charters, et al. Guidelines for performing systematic literature reviews in software engineering version 2.3. *Engineering*, 45(4ve):1051, 2007.
- [121] F. Klauser. Smart heritage: Interoperability and ethics in digital cultural systems. *Heritage Science*, 9(1):42, 2021.

- [122] M. Koeva, M. Luleva, and P. Maldjanski. Integrating spherical panoramas and maps for visualization of cultural heritage objects using virtual reality technology. *Sensors*, 17(4):829, 2017.
- [123] M. N. Koeva. 3d modelling and interactive web-based visualization of cultural heritage objects. *The International Archives of the Photogrammetry, Remote Sensing and Spatial Information Sciences*, 41:297–303, 2016.
- [124] X. Kong. Monitoring time-varying changes of historic structures through photogrammetry-driven digital twinning. *arXiv preprint arXiv:2407.18925*, 2024.
- [125] M. Konstantakis, J. Aliprantis, A. Teneketzi, and G. Caridakis. Understanding user experience aspects in cultural heritage interaction. In *Proceedings of the 22nd Pan-Hellenic Conference on Informatics*, pages 267–271, 2018.
- [126] M. Konstantakis, G. Trichopoulos, J. Aliprantis, N. Gavogiannis, A. Karagianni, P. Parthenios, K. Serraios, and G. Caridakis. An improved approach for generating digital twins of cultural spaces through the integration of photogrammetry and laser scanning technologies. *Digital*, 4(1):215–231, 2024.
- [127] H. Kopetz. *Real-Time Systems: Design Principles for Distributed Embedded Applications*. Springer, 2011.
- [128] M. Kozlovskikh, I. Neverova, and N. Ustinova. Virtual reality tools for creating interactive digital twins of attractions infrastructure. In G. Kazeeva, A. Kolyada, and S. Berseneva, editors, *Immersive Technologies in Virtual Heritage for an Innovative Tourist Experience: The Case of the Berati Ethnographic Museum*, pages 123–135. Springer, 2024.
- [129] W. Kritzinger, M. Karner, G. Traar, J. Henjes, and W. Sihn. Digital twin in manufacturing: A categorical literature review and classification. *Ifac-PapersOnline*, 51(11):1016–1022, 2018.
- [130] T. Kuflik, Z. Boger, and M. Zancanaro. Analysis and prediction of museum visitors’ behavioral pattern types. *Ubiquitous Display Environments*, pages 161–176, 2012.
- [131] Z. Kvítková and Z. Petru. Cultural unesco heritage in covid-19 pandemic times. *Academica Turistica-Tourism and Innovation Journal*, 16(1), 2023.
- [132] F. La Russa and C. Santagati. Historical sentient–building information model: a digital twin for the management of museum collections in historical architectures. *The international archives of the photogrammetry, remote sensing and spatial information sciences*, 43:755–762, 2020.
- [133] F. Landi, M. G. Bevilacqua, G. Caroti, F. Meligeni, A. Piemonte, P. Rechichi, and P. Croce. Digital twin for structural health monitoring of cultural heritage: The buildchain demo-pilot, palazzo poniatowski-guadagni in florence. In *International Conference on Communication and Computational Technologies*, pages 191–204. Springer, 2024.
- [134] N. Laohaviraphap and T. Waroonkun. Integrating artificial intelligence and the internet of things in cultural heritage preservation: A systematic review of risk management and environmental monitoring strategies. *Buildings*, 14(12):3979, 2024.

- [135] F. Larue, M. Di Benedetto, M. Dellepiane, and R. Scopigno. From the digitization of cultural artifacts to the web publishing of digital 3d collections: an automatic pipeline for knowledge sharing. *Journal of multimedia*, 7(2), 2012.
- [136] C. Lee, H. Kwon, and Y. I. Lee. Opc ua-based three-layer architecture for aggregated microgrids integrating edge cloud computing and iec 62264. *Journal of Industrial Information Integration*, page 100965, 2025.
- [137] L. Li, S. Aslam, A. Wileman, and S. Perinpanayagam. Digital twin in aerospace industry: A gentle introduction. *IEEE Access*, 10:9543–9562, 2021.
- [138] M. Li, S. Xu, J. Tang, and W. Chen. Design and research of digital twin platform for handicraft intangible cultural heritage—yangxin cloth paste. *Heritage Science*, 12(1), 2024.
- [139] T. Li, A. K. Sahu, M. Zaheer, M. Sanjabi, and H. B. McMahan. Federated learning: Challenges, methods, and future directions. *IEEE Signal Processing Magazine*, 37(3):50–60, 2020.
- [140] X. Li, Y. Zhang, X. Zhang, and Z. Li. Digital twin of wooden heritage through multidimensional model construction and integration. *Structural Control Health Monitoring*, 31(1):e2696, 2024.
- [141] Y. Lian and J. Xie. *The evolution of digital cultural heritage research: Identifying key trends, hotspots, and challenges through bibliometric analysis*, volume 16. MDPI, 2024.
- [142] F. V. Liddell. *The Crypto-Museum: Investigating the impact of blockchain and NFTs on digital ownership, authority, and authenticity in museums*. The University of Manchester (United Kingdom), 2022.
- [143] K. Y. H. Lim, P. Zheng, and C.-H. Chen. A state-of-the-art survey of digital twin: techniques, engineering product lifecycle management and business innovation perspectives. *Journal of Intelligent Manufacturing*, 31(6):1313–1337, 2020.
- [144] W. Y. B. Lim et al. Federated learning for smart cities: Privacy-preserving urban computing. *IEEE Internet of Things Journal*, 7(10):10745–10755, 2020.
- [145] C. L. Liu and J. W. Layland. Scheduling algorithms for multiprogramming in a hard-real-time environment. *Journal of the ACM (JACM)*, 20(1):46–61, 1973.
- [146] L. Liu, H. Jin, Y. Liu, and X. Zhang. Intelligent evacuation route planning algorithm based on maximum flow. *International journal of environmental research and public health*, 19(13):7865, 2022.
- [147] W. Liu, M. Lu, Y. Chen, and K. Yan. Digital twin applications in an archaeological site: A virtual reconstruction of the pishan site, zhejiang, china. *Journal of Chinese Architecture and Urbanism*, 6(1):1735, 2024.
- [148] X. Liu, D. Jiang, B. Tao, F. Xiang, G. Jiang, Y. Sun, J. Kong, and G. Li. A systematic review of digital twin about physical entities, virtual models, twin data, and applications. *Advanced Engineering Informatics*, 55:101876, 2023.

- [149] X. Liu and J. Wang. Digital twin-based smart museum: A framework for artifact preservation and visitor experience. *Sensors*, 20(24):7192, 2020.
- [150] Z. Liu and S. Chang. A study of digital exhibition visual design led by digital twin and vr technology. *Measurement: Sensors*, 31, 2024.
- [151] Z. Liu, J. Guo, W. Yang, J. Fan, K.-Y. Lam, and J. Zhao. Privacy-preserving aggregation in federated learning: A survey. *IEEE Transactions on Big Data*, 2022.
- [152] F. J. López, P. M. Leronés, J. Llamas, J. Gómez-García-Bermejo, and E. Zalama. A review of heritage building information modeling (h-bim). *Multimodal technologies and interaction*, 2(2):21, 2018.
- [153] R. Lorenzo, B. Fabrice, D.-V. Paolo, D. Jean-Christophe, G. Giacomo, H. Franck, J. Delphine, M. Paola, T. Marco, R. Elisabeth, et al. Learning from objects: the use of advanced numerical methods to exploit a complete set of information from experimental data, for the mona lisa’s digital-twin. In *IOP Conference Series: Materials Science and Engineering*, volume 949, page 012089. IOP Publishing, 2020.
- [154] Y. Lu. From digital twin to intelligent twin: A roadmap for smart manufacturing. *Computers in Industry*, 128:103440, 2022.
- [155] Y. Lu, X. Huang, K. Zhang, S. Maharjan, and Y. Zhang. Communication-efficient federated learning for digital twin edge networks in industrial iot. *IEEE Transactions on Industrial Informatics*, 17(8):5709–5718, 2020.
- [156] Y. Lu, X. Xu, and L. Wang. Digital twins in smart grids: State-of-the-art, challenges, and opportunities. *Renewable and Sustainable Energy Reviews*, 133:110157, 2020.
- [157] Y. Lu, Z. Yang, X. Zheng, J. Wang, and D. Kiritsis. Exploring the concept of cognitive digital twin from model-based systems engineering perspective. *The International Journal of Advanced Manufacturing Technology*, 121:5835–5854, 2022.
- [158] E. Lucchi. Digital twins for the automation of the heritage construction sector. *Automation in Construction*, 156:105073, 2023.
- [159] W. Luo, Z. Yu, R. Rzayev, M. Satkowski, S. Gumhold, M. McGinity, and R. Dachsel. Pearl: Physical environment based augmented reality lenses for in-situ human movement analysis. In *Proceedings of the 2023 CHI Conference on Human Factors in Computing Systems*, pages 1–15, 2023.
- [160] W. Luther, N. Baloian, D. Biella, and D. Sacher. Digital twins and enabling technologies in museums and cultural heritage: An overview. *Sensors*, 23(3):1583, 2023.
- [161] R. Maalek and S. Maalek. Automatic recognition and digital documentation of cultural heritage hemispherical domes using images. *ACM Journal on Computing and Cultural Heritage*, 16(1):1–21, 2022.
- [162] A. Madni, C. Madni, and S. Lucero. Architecting cyber–physical systems of systems. *Systems*, 7(1):1–20, 2019.

- [163] H. Marah and M. Challenger. An architecture for intelligent agent-based digital twin for cyber-physical systems. In *Digital twin driven intelligent systems and emerging metaverse*, pages 65–99. Springer, 2023.
- [164] H. Marah and M. Challenger. Adaptive hybrid reasoning for agent-based digital twins of distributed multi-robot systems. *Simulation*, 100(9):931–957, 2024.
- [165] R. Markov, G. Bogdanova, and M. Ivanova. Creating and storing 7d digital twins. In *2023 8th International Conference on Smart and Sustainable Technologies (SpliTech)*, pages 1–5. IEEE, 2023.
- [166] A. Marra, S. Gerbino, A. Greco, and G. Fabbrocino. Combining integrated informative system and historical digital twin for maintenance and preservation of artistic assets. *Sensors*, 21(17):5956, 2021.
- [167] A. Massafra, G. Predari, R. Gulli, et al. Towards digital twin driven cultural heritage management: A hbim-based workflow for energy improvement of modern buildings. *International Archives of the Photogrammetry, Remote Sensing and Spatial Information Sciences*, 46:149–157, 2022.
- [168] M. Mattila. Government resolution for the cultural heritage strategy 2023–2030. 2023.
- [169] S. Mazzetto. Integrating emerging technologies with digital twins for heritage building conservation: an interdisciplinary approach with expert insights and bibliometric analysis. *Heritage*, 7(11):6432–6479, 2024.
- [170] S. Mazzetto. Integrating emerging technologies with digital twins for heritage building conservation: An interdisciplinary approach with expert insights and bibliometric analysis. *Heritage*, 7(11):6432–6479, 2024.
- [171] J. McCaffery, A. Miller, A. Vermehren, and A. Fabola. The virtual museums of caen: A case study on modes of representation of digital historical content. In *2015 Digital Heritage*, volume 2, pages 541–548. IEEE, 2015.
- [172] H. B. McMahan, E. Moore, D. Ramage, S. Hampson, and B. A. Arcas. Communication-efficient learning of deep networks from decentralized data. In *AISTATS 2017*, pages 1273–1282, 2017.
- [173] S. Mehta, V. Kukreja, and R. Gupta. Exploring the potential of federated learning cnn for interactive virtual tours of unesco cultural heritage sites: A case study. In *2023 14th International Conference on Computing Communication and Networking Technologies (ICCCNT)*, pages 1–6. IEEE, 2023.
- [174] O. Menaguale. Digital twin and cultural heritage—the future of society built on history and art. In *The Digital Twin*, pages 1081–1111. Springer, 2023.
- [175] D. Menon, B. Anand, and C. L. Chowdhary. Digital twin: exploring the intersection of virtual and physical worlds. *IEEE Access*, 2023.

- [176] B. Metcalfe, H. C. Boshuizen, J. Bulens, and J. J. Koehorst. Digital twin maturity levels: a theoretical framework for defining capabilities and goals in the life and environmental sciences. *F1000Research*, 12:961, 2023.
- [177] V. Mittal, M. Jahanian, and K. Ramakrishnan. Flare: Federated active learning assisted by naming for responding to emergencies. In *Proceedings of the 8th ACM Conference on Information-Centric Networking*, pages 71–82, 2021.
- [178] J. Moyano, E. Cabrera-Revuelta, J. E. Nieto-Julián, M. Fernández-Alconchel, and P. Fernández-Valderrama. Evaluation of geometric data registration of small objects from non-invasive techniques: Applicability to the hbim field. *Sensors*, 23(3):1730, 2023.
- [179] S. S. Mullins and B. F. G. Katz. The past has ears at notre-dame: Acoustic digital twins for research and narration. *Acoustics Today*, 20(2):49, 2024.
- [180] I. Munoz-Pandiella, C. Bosch, M. Guardia, B. Cayuela, P. Pogliani, G. Bordi, M. Paschali, C. Andújar, and P. Charalambous. Digital twins for medieval monuments: Requirements from art historians and technical challenges for analysis and restoration. In *2022 International Conference on Interactive Media, Smart Systems and Emerging Technologies (IMET)*, pages 1–8. IEEE, 2022.
- [181] M. Murphy, S. Fai, L. Chow, E. Meegan, S. Scandurra, S. Pavia, A. Corns, and J. Cahil. Virtual historic centers: Digital representation of archaeological heritage. In *Handbook of Cultural Heritage Analysis*, pages 1497–1528. Springer, 2022.
- [182] J.-H. Nancarrow. Countering the "digital uncanny": Post-processing for 3-d digital heritage. *Studies in Digital Heritage*, 3(2):170–185, 2019.
- [183] E. Negri, L. Fumagalli, and M. Macchi. A review of the roles of digital twin in cps-based production systems. *Procedia Manufacturing*, 11:939–948, 2017.
- [184] D. T. Ngo, C. V. Ta, C. T. Ma, H. M. Nguyen, H. X. Nguyen, and H. T. Le. Quality assessment criteria and methods for 3d digital replica of historical printing woodblocks. *ACM Journal on Computing and Cultural Heritage*, 16(3):1–16, 2023.
- [185] F. Niccolucci, A. Felicetti, and S. Hermon. Populating the data space for cultural heritage with heritage digital twins. *Data*, 7(8):105, 2022.
- [186] C. Nicholson, S. Kansa, N. Gupta, and R. Fernandez. Will it ever be fair?: Making archaeological data findable, accessible, interoperable, and reusable. *Advances in Archaeological Practice*, 11(1):63–75, 2023.
- [187] E. Nieto-Julián, M. D. Robador, J. Moyano, and S. Bruno. Semantic hbim for heritage conservation: A methodology for mapping deterioration and structural deformation in historic envelopes. *Buildings*, 15(12):1990, 2025.
- [188] W. Ning. Innovative frameworks for protecting intangible cultural heritage using advanced multimedia technology drawing on digital twins and 3d modeling. *Intelligent Decision Technologies*, 19(2):1023–1039, 2024.

- [189] E. Nofal, M. Alshammari, M. Alghamdi, and M. Alshahrani. Risk management in cultural heritage sites: Toward digital preparedness. *Sustainability*, 13(5):2540, 2021.
- [190] R. Ogunsakin, N. Mehandjiev, and C. A. Marin. Towards adaptive digital twins architecture. *Computers in Industry*, 149:103920, 2023.
- [191] D.-G. J. Opoku, S. Perera, R. Osei-Kyei, M. Rashidi, T. Famakinwa, and K. Bamdad. Drivers for digital twin adoption in the construction industry: A systematic literature review. *Buildings*, 12(2):113, 2022.
- [192] F. ÖZER SARI and M. NAZLI. Sustaining cultural heritage by means of museums in an ever-changing world. *Gaziantep Üniversitesi Sosyal Bilimler Dergisi*, 17(1), 2018.
- [193] X. Pan, Q. Lin, S. Ye, L. Li, L. Guo, and B. Harmon. Deep learning based approaches from semantic point clouds to semantic bim models for heritage digital twin. *Heritage Science*, 12(1):65, 2024.
- [194] U. Paris and F. Mairesse. *Museums Around the World in the Face of COVID-19*. PhD thesis, Unesco, 2021.
- [195] P. Pei, Z. Chen, and Y. Wang. Production analytics for renewable energy portfolio management: A comparative study of active and passive strategies in digital twin. *Solar Energy*, 262:111923, 2023.
- [196] F. Piccialli, S. Cuomo, V. S. d. Cola, and G. Casolla. A machine learning approach for iot cultural data. *Journal of Ambient Intelligence and Humanized Computing*, pages 1–12, 2019.
- [197] E. Pietroni and D. Ferdani. Virtual restoration and virtual reconstruction in cultural heritage: Terminology, methodologies, visual representation techniques and cognitive models. *Information*, 12(4):167, 2021.
- [198] F. S. Pomposo, D. O. González, H. F. Aparicio, I. J. I. Berrocal, and A. Naddeo. Creation of the digital twin of the cabo de palos lighthouse by reverse engineering. In *International conference on The Digital Transformation in the Graphic Engineering*, pages 119–129. Springer, 2023.
- [199] F. Poux, R. Neuville, P. Hallot, L. Van Wersch, A. Luczfalvy Jancsó, and R. Billen. Digital investigations of an archaeological smart point cloud: A real time web-based platform to manage the visualisation of semantical queries. In *GEOMATICS & RESTORATION 2017*, 2017.
- [200] E. Puertas, F. Ávila, E. García-Macías, and R. Gallego. Preventive preservation of rammed earth historical heritage through continuous monitoring, architectural inspections, and data fusion. *Buildings*, 14(10):3294, 2024.
- [201] F. Qamar, M. Rizwan, and S. Kiani. A simulation model of federated learning in multi-agent systems using netlogo. *Simulation Modelling Practice and Theory*, 120:102520, 2022.
- [202] Q. Qi and F. Tao. Digital twin for smart manufacturing: A review. *Journal of Manufacturing Systems*, 48:308–321, 2018.

- [203] Y. Qu, L. Gao, Y. Xiang, S. Shen, and S. Yu. Fedtwin: Blockchain-enabled adaptive asynchronous federated learning for digital twin networks. *IEEE Network*, 36(6):183–190, 2022.
- [204] R. Rai and C. K. Sahu. Driven by data or derived through physics? a review of hybrid physics guided machine learning techniques with cyber-physical system (cps) focus. *IEEE Access*, 8:71050–71073, 2020.
- [205] S. Railsback and V. Grimm. *Agent-Based and Individual-Based Modeling: A Practical Introduction*. Princeton University Press, 2012.
- [206] S. F. Railsback and V. Grimm. *Agent-based and individual-based modeling: a practical introduction*. Princeton university press, 2019.
- [207] 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*, 79:103663, 2022.
- [208] A. Rasheed, O. San, and T. Kvamsdal. Digital twin: Values, challenges and enablers. *arXiv preprint arXiv:1910.01719*, 2019.
- [209] M. Redeker, J. N. Weskamp, B. Rössl, and F. Pethig. Towards a digital twin platform for industrie 4.0. In *2021 4th IEEE international conference on industrial cyber-physical systems (ICPS)*, pages 39–46. IEEE, 2021.
- [210] I. Rocca, L. Berto, G. D’Acunto, and A. Saetta. Seismic vulnerability assessment of venetian façades and artistic assets: Historical digital twin approach. *Journal of Building Engineering*, 95:109992, 2024.
- [211] E. J. Sacoto-Cabrera, A. Perez-Torres, L. Tello-Oquendo, and M. Cerrada. Iot, ai, and digital twins in smart cities: A systematic review for a thematic mapping and research agenda. *Smart Cities*, 8(5):175, 2025.
- [212] M. Safari, R. Afkhami, M. Amerzadeh, and V. Zaroushani. Simulation of fire emergency evacuation in a primary school based on pathfinder software. *Buildings*, 15(1):90, 2024.
- [213] M. Schranz, G. A. Di Caro, T. Schmickl, W. Elmenreich, F. Arvin, A. Şekercioglu, and M. Sende. Swarm intelligence and cyber-physical systems: concepts, challenges and future trends. *Swarm and Evolutionary Computation*, 60:100762, 2021.
- [214] J. Schulman, F. Wolski, P. Dhariwal, A. Radford, and O. Klimov. Proximal policy optimization algorithms. *arXiv preprint arXiv:1707.06347*, 2017.
- [215] A. Scianna, G. Gaglio, and M. La Guardia. Digital photogrammetry, tls survey and 3d modelling for vr and ar applications in ch. *The International Archives of the Photogrammetry, Remote Sensing and Spatial Information Sciences*, 43:901–909, 2020.
- [216] G. Selim, M. Jamhawi, M. G. Abdelmonem, S. Ma’bdeh, and A. Holland. The virtual living museum: Integrating the multi-layered histories and cultural practices of gadara’s archaeology in umm qais, jordan. *Sustainability*, 14(11):6721, 2022.

- [217] A. Shabani, M. Skamantzari, S. Tapinaki, A. Georgopoulos, V. Plevris, and M. Kioumarsis. 3d simulation models for developing digital twins of heritage structures: challenges and strategies. *Procedia structural integrity*, 37:314–320, 2022.
- [218] N. Shahzad, W. de Paula Ferreira, and F. Deschamps. Cognitive digital twins: A state-of-the-art review. *Available at SSRN 5085883*.
- [219] S. Shang. Exploring the path of digital governance of urban industrial pollution: empirical evidence from 280 cities in china. *Environmental Science and Pollution Research*, 31(43):55388–55409, 2024.
- [220] X. Shang, M. Guo, G. Wang, J. Zhao, and D. Pan. Behavioral model construction of architectural heritage for digital twin. *npj Heritage Science*, 13, 2025.
- [221] A. Sharma, E. Kosasih, J. Zhang, A. Brintrup, and A. Calinescu. Digital twins: State of the art theory and practice, challenges, and open research questions. *Journal of Industrial Information Integration*, page 100383, 2022.
- [222] R. Sitnik and W. Zaluski. 3dmac–3d digitization of european cultural heritage. In *Proceedings of the Symposium on Photonics Technologies for 7th Framework Program*, volume 502. Citeseer, 2006.
- [223] V. Smith, C. Chiang, M. Sanjabi, and A. Talwalkar. Federated multi-task learning. In *Advances in Neural Information Processing Systems (NeurIPS) 30*, 2017.
- [224] F. Stanco and D. Tanasi. Beyond virtual replicas: 3d modeling and maltese prehistoric architecture. *Journal of Electrical and Computer engineering*, 2013:11–11, 2013.
- [225] J. A. Stankovic, M. Spuri, M. Di Natale, and G. C. Buttazzo. Implications of classical scheduling results for real-time systems. *Computer*, 28(6):16–25, 1995.
- [226] M. Štroner, T. Křemen, and R. Urban. Progressive dilution of point clouds considering the local relief for creation and storage of digital twins of cultural heritage. *Applied Sciences*, 12(22):11540, 2022.
- [227] E. Stylianidis. Cipa-heritage documentation: 50 years: looking backwards. *The International Archives of Photogrammetry, Remote Sensing and Spatial Information Sciences*, 42:1–130, 2019.
- [228] E. Stylianidis and F. Remondino. *3D Recording, Documentation and Management of Cultural Heritage*. Whittles Publishing, 2016.
- [229] Y. Su, Z. Zhou, L. Zhao, and Y. Wang. China’s exploration and experience in digital twinning of tangible cultural heritage. In *Intelligent Computing for Cultural Heritage*. Routledge, 2024.
- [230] G. Sugiyama, I. Bourgeois, and H. Rodrigues. A holistic methodology for the assessment of heritage digital twin applied to portuguese case studies. *Digital Applications in Archaeology and Cultural Heritage*, 36:e00390, 2025.

- [231] T. Sun, D. Li, and B. Wang. Decentralized federated averaging. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 45(4):4289–4301, 2022.
- [232] W. Sun, S. Lei, L. Wang, Z. Liu, and Y. Zhang. Adaptive federated learning and digital twin for industrial internet of things. *IEEE Transactions on Industrial Informatics*, 17(8):5605–5614, 2020.
- [233] Y. Sun, Z. Peng, H. Zhou, et al. Adaptive federated learning in edge computing for industrial cyber-physical systems. *IEEE Transactions on Industrial Informatics*, 17(3):1935–1945, 2021.
- [234] E. Suparto and D. Erwandi. Literature review: Human behaviour and evacuation fire system. *Asian Journal of Engineering, Social and Health*, 3(6):1284–1299, 2024.
- [235] N. Suri, A. Gokhale, and H. Neema. Real-time embedded systems: Simulation, modeling and analysis. *ACM Transactions on Embedded Computing Systems*, 15(4):1–29, 2016.
- [236] J. Tan, J. Leng, X. Zeng, D. Feng, and P. Yu. Digital twin for xiegong’s architectural archaeological research: A case study of xuanluo hall, sichuan, china. *Buildings*, 12(7):1053, 2022.
- [237] X. Tang, Q. Zhong, and X. Zheng. Exploring the design of digital twin virtual space for the awaken lion culture based on technology acceptance modeling. In *Proceedings of the Eleventh International Symposium of Chinese CHI*, pages 225–240, 2023.
- [238] F. Tao and Q. Qi. Digital twin-driven intelligent systems for self-learning factories. *Advanced Engineering Informatics*, 52:101553, 2022.
- [239] F. Tao, H. Zhang, A. Liu, and A. Y. Nee. Digital twin in industry: State-of-the-art. *IEEE Transactions on industrial informatics*, 15(4):2405–2415, 2018.
- [240] F. Tao and M. Zhang. Digital twins and cyber-physical systems: A new paradigm for future smart industries. *IEEE Access*, 7:97270–97280, 2019.
- [241] F. Tao, M. Zhang, Y. Liu, and A. Y. C. Nee. Digital twin in industry: State-of-the-art. *IEEE Access*, 6:3585–3593, 2018.
- [242] S. Taylor, A. Q. Rodriguez, and L. Scudiero. D2. 01: Initial guidelines on data management. 2018.
- [243] T. Teba, H. Hawash, A. H. Moustafa, and N. Gharib. A tailored framework for heritage and values identification in historic alexandria: Cultural mapping through historic urban landscape conceptualisation. *Journal of Cultural Heritage Management and Sustainable Development*, 2025.
- [244] K. Themistocleous. The monitoring of cultural heritage sites using digital twins. In K. Schulz, K. G. Nikolakopoulos, and U. Michel, editors, *Earth Resources and Environmental Remote Sensing/GIS Applications XV*. SPIE, Nov. 2024.
- [245] J. Thomas and A. Harden. Methods for the thematic synthesis of qualitative research in systematic reviews. *BMC medical research methodology*, 8:1–10, 2008.

- [246] S. Tisue and U. Wilensky. Netlogo: A simple environment for modeling complexity. In *International Conference on Complex Systems*, 2004.
- [247] V. Titarenko and I. Kaliukh. Digital twins applying to the assessment of the hostynnyi dvir historical monument destruction. In *International Conference on Protection of Historical Constructions*, pages 569–576. Springer, 2025.
- [248] F. Tschirschwitz, G. Büyüksalih, T. Kersten, T. Kan, G. Enc, and P. Baskaraca. Virtualising an ottoman fortress—laser scanning and 3d modelling for the development of an interactive, immersive virtual reality application. *The International Archives of the Photogrammetry, Remote Sensing and Spatial Information Sciences*, 42:723–729, 2019.
- [249] J.-F. Uhlenkamp, J. B. Hauge, E. Broda, M. Lütjen, M. Freitag, and K.-D. Thoben. Digital twins: A maturity model for their classification and evaluation. *IEEE Access*, 10:69605–69635, 2022.
- [250] UNESCO. Recommendation on the ethics of artificial intelligence. Technical report, UNESCO, 2021.
- [251] UNESCO. Recommendation on the ethics of artificial intelligence, 2023.
- [252] United Nations. Sustainable development goals report 2020. United Nations Publishing, 2020.
- [253] URBACT. Mostar’s metacity: Pioneering a digital and sustainable future, 2025. Accessed: 2025-08-30.
- [254] M. Uysal and Others. Detection of cracks in cultural heritage buildings using uav photogrammetry-based digital twin. *Journal of Building Engineering*, 47:103289, 2024.
- [255] F. Valeonti, A. Bikakis, M. Terras, C. Speed, A. Hudson-Smith, and K. Chalkias. Crypto collectibles, museum funding and.openglam: challenges, opportunities and the potential of non-fungible tokens (nfts). *Applied Sciences*, 11(21):9931, 2021.
- [256] E. Vila-Chã, A. Barontini, and P. B. Lourenço. Implementation of a condition monitoring strategy for the monastery of salzedas, portugal: Challenges and optimisation. *Buildings*, 13(3):719, 2023.
- [257] ViMM. Vimm - virtual multimodal museums, 2025. Accessed: 2025-08-30.
- [258] A. Vuoto, M. F. Funari, and P. B. Lourenço. Shaping digital twin concept for built cultural heritage conservation: A systematic literature review. *International journal of architectural heritage*, pages 1–34, 2023.
- [259] D. Wagg, K. Worden, R. Barthorpe, and P. Gardner. Digital twins: state-of-the-art and future directions for modeling and simulation in engineering dynamics applications. *ASCE-ASME Journal of Risk and Uncertainty in Engineering Systems, Part B: Mechanical Engineering*, 6(3):030901, 2020.
- [260] F. Wang and J. Du. Digital presentation and communication of cultural heritage in the post-pandemic era. *ISPRS Annals of the Photogrammetry, Remote Sensing and Spatial Information Sciences*, 8:187–192, 2021.

- [261] J. Wang, Y. Li, R. X. Gao, and F. Zhang. Hybrid physics-based and data-driven models for smart manufacturing: Modelling, simulation, and explainability. *Journal of Manufacturing Systems*, 63:381–391, 2022.
- [262] M. Wang and N. Lau. Nft digital twins: A digitalization strategy to preserve and sustain miao silver craftsmanship in the metaverse era. *Heritage*, 6(2):1921–1941, 2023.
- [263] P. Wang, X. Ma, L. Fei, H. Zhang, D. Zhao, and J. Zhao. When the digital twin meets the preventive conservation of movable wooden artifacts. *Heritage Science*, 11(1):1–20, 2023.
- [264] R. Wang, M. Guo, Y. Zhang, J. Chen, Y. Wei, and L. Zhu. Three-dimensional intelligent understanding and preventive conservation prediction for linear cultural heritage. *Buildings*, 15(16):2827, 2025.
- [265] S. Wang, L. Li, X. Zhang, and J. Li. Optimization of structural reinforcement assessment for architectural heritage digital twins based on lidar and multi-source remote sensing. *Heritage Science*, 12(1):404, 2024.
- [266] Y. Wang, J. Ge, and A. Comber. An evacuation simulation model of pedestrian flow using bayesian nash equilibrium and a multi-agent system. *AGILE: GIScience Series*, 3:68, 2022.
- [267] Y. Wang, M. Kyriakidis, and V. N. Dang. Incorporating human factors in emergency evacuation—an overview of behavioral factors and models. *International Journal of Disaster Risk Reduction*, 60:102254, 2021.
- [268] Y. Wang, Z. Su, S. Guo, M. Dai, T. H. Luan, and Y. Liu. A survey on digital twins: Architecture, enabling technologies, security and privacy, and future prospects. *IEEE Internet of Things Journal*, 10(17):14965–14987, 2023.
- [269] S. Wessels, S. Maape, B. J. Schoville, and J. Wilkins. The drone, the snake, and the crystal: Manifesting potency in 3d digital replicas of living heritage and archaeological places. *Archaeologies*, 19(2):395–429, 2023.
- [270] M. A. Wijnhoven and A. Moskvin. Digital replication and reconstruction of mail armour. *Journal of Cultural Heritage*, 45:221–233, 2020.
- [271] U. Wilensky. *NetLogo*. Center for Connected Learning and Computer-Based Modeling, Northwestern University, 1999.
- [272] R. Wilhelm, J. Engblom, A. Ermedahl, et al. The worst-case execution-time problem—overview of methods and survey of tools. *ACM Transactions on Embedded Computing Systems*, 7(3):1–53, 2008.
- [273] F. Windhager, P. Federico, G. Schreder, K. Glinka, M. Dörk, S. Miksch, and E. Mayr. Visualization of cultural heritage collection data: State of the art and future challenges. *IEEE transactions on visualization and computer graphics*, 25(6):2311–2330, 2018.
- [274] C. Wohlin. Guidelines for snowballing in systematic literature studies and a replication in software engineering. In *Proceedings of the 18th international conference on evaluation and assessment in software engineering*, pages 1–10, 2014.

- [275] K. Worden, E. Cross, R. Barthorpe, D. Wagg, and P. Gardner. On digital twins, mirrors, and virtualizations: Frameworks for model verification and validation. *ASCE-ASME Journal of Risk and Uncertainty in Engineering Systems, Part B: Mechanical Engineering*, 6(3):030902, 2020.
- [276] R. Wu, L. Gao, H. Lee, J. Xu, and Y. Pan. A study of the key factors influencing young users' continued use of the digital twin-enhanced metaverse museum. *Electronics*, 13(12):2303, 2024.
- [277] L. Xia, L. Zhang, and J. Zhang. Autonomous digital twin for real-time decision-making in cyber-physical systems. *Computers in Industry*, 146:103844, 2023.
- [278] L. Xu, S. Yan, Z. Chen, and X. Chen. Design of the museum interactive lighting system based on the digital twin technology. *Scientific Programming*, 2021:1–8, 2021.
- [279] S. Xu and H. Chen. Deep learning and digital twin integration for structural damage detection in ancient pagodas. *Scientific Reports*, 15(1):28408, 2025.
- [280] Q. Yang, Y. Liu, T. Chen, and Y. Tong. Federated machine learning: Concept and applications. *ACM Transactions on Intelligent Systems and Technology*, 10(2):1–19, 2019.
- [281] A. Y. Yiğit and M. Uysal. Automatic crack detection and structural inspection of cultural heritage buildings using uav photogrammetry and digital twin technology. *Journal of Building Engineering*, 94:109952, 2024.
- [282] A. Ynnerman, T. Rydell, D. Antoine, D. Hughes, A. Persson, and P. Ljung. Interactive visualization of 3d scanned mummies at public venues. *Communications of the ACM*, 59(12):72–81, 2016.
- [283] C. Zhang, S. Guo, J. Wu, et al. Asynchronous federated learning for smart industrial internet of things. *IEEE Transactions on Industrial Informatics*, 18(4):2547–2557, 2022.
- [284] W. Zhang and N. Taib. Virtual experimental model and digital preservation path of traditional wooden carvings in huizhou traditional dwellings. *International Journal of Conservation Science*, 16(1):199–206, 2025.
- [285] X. Zhang, D. Yang, C. H. Yow, L. Huang, X. Wu, X. Huang, J. Guo, S. Zhou, and Y. Cai. Metaverse for cultural heritages. *Electronics*, 11(22):3730, 2022.
- [286] Y. Zhang, M. Chen, and D. Wu. Federated learning for intelligent transportation systems: Recent advances and future challenges. *IEEE Internet of Things Journal*, 8(8):7092–7105, 2021.
- [287] Z. Zhang, A. Dang, and J. Liu. Advancing conservation methods of the great wall cultural heritage through digital twin. *IEEE Internet Computing*, 29(1):20–30, 2025.
- [288] J. Zhao, L. Guo, and Y. Li. Application of digital twin combined with artificial intelligence and 5g technology in the art design of digital museums. *Wireless Communications and Mobile Computing*, 2022, 2022.

- [289] L. Zhao, L. Wang, F. Yu, Y. Zhao, and Y. Wang. Will users prefer cultural heritage exhibition in digital twin? the mediating role of perceived authenticity under the moderation of individual profiles. *Proceedings of the Association for Information Science and Technology*, 61(1):709–714, 2024.
- [290] Z. Zheng, Y. Zhou, Y. Sun, Z. Wang, B. Liu, and K. Li. Applications of federated learning in smart cities: recent advances, taxonomy, and open challenges. *Connection Science*, 34(1):1–28, 2022.
- [291] B. Zhou, M. Yan, and G. Xia. Study on digital protection path of archaeological ruins of liangzhu city. In *Proceedings of the 2024 5th International Conference on Computer Science, Electronics and Communication Engineering (CSECE)*, pages 1–7, 2024.
- [292] L. Zhou, S. Leng, and Q. Wang. A federated digital twin framework for uavs-based mobile scenarios. *IEEE Transactions on Mobile Computing*, 23(6):7377–7393, 2023.
- [293] L. Zhou, S. Yang, X. Zhang, Y. Liu, and X. Xu. Ai-powered digital twins: Intelligent maintenance and operation. *Advanced Engineering Informatics*, 50:101395, 2021.
- [294] Y. Zhou, Q. Ye, and J. Lv. Communication-efficient federated learning with compensated overlap-fedavg. *IEEE Transactions on Parallel and Distributed Systems*, 33(1):192–205, 2021.
- [295] C. Zou, S.-Y. Rhee, L. He, D. Chen, and X. Yang. Sounds of history: A digital twin approach to musical heritage preservation in virtual museums. *Electronics*, 13(12):2388, 2024.
- [296] B. Świt Jankowska. Adam mickiewicz museum in Śmiełów – towards a contemporary museum concept using digital twin technology. *Architectus*, 79(3):65–74, 2024.

Appendix: Mathematical Validation of the Federated Digital Twin Framework

This section provides a mathematical validation of the proposed Federated Digital Twin (FDT) architecture visualized in Figure 4.2. The validation focuses on the core components of local digital twins, the federated learning process, the aggregation mechanism, and emergency decision-making logic.

1. Local Digital Twin Modeling

Each museum $i \in \{1, 2, \dots, N\}$ has a local digital twin characterized by:

- Sensor data: $S_i(t)$
- Local evacuation decision function: $E_i(t)$
- Data analytics process: $A_i(t)$

The evacuation decision is computed as:

$$E_i(t) = f_i(S_i(t), A_i(t)) \Rightarrow R_i(t)$$

where $R_i(t)$ is the optimal evacuation route recommendation for museum i at time t .

2. Federated Learning and Model Aggregation

Each local twin trains a model using its data and contributes only the updated model parameters, not raw data, preserving privacy. The training step for each local model is:

$$w_i^{(t+1)} = w_i^{(t)} - \eta \nabla \mathcal{L}_i(w_i^{(t)})$$

The global model is updated using a weighted average of all local models:

$$w_{\text{global}}^{(t+1)} = \sum_{i=1}^N \frac{n_i}{n} w_i^{(t+1)}$$

where:

- $\mathcal{L}_i$: Local loss function
- w_i : Model parameters at site i

- η : Learning rate
- n_i : Number of samples at museum i
- $n = \sum_{i=1}^N n_i$: Total number of samples

3. Emergency Simulation and Decision Constraints

The safety and evacuation efficiency are constrained by:

$$\text{Evacuation_Success}(i, t) = \begin{cases} 1, & \text{if } T_i^{\text{evac}}(t) \leq T_{\text{safe}} \\ 0, & \text{otherwise} \end{cases}$$

where T_i^{evac} is the evacuation time, and T_{safe} is the safety threshold.

4. Security and Privacy Layer

The federated learning protocol integrates privacy-preserving mechanisms:

- Differential Privacy: $\tilde{w}_i^{(t)} = w_i^{(t)} + \mathcal{N}(0, \sigma^2)$
- Encryption Function: $C_i = \text{Enc}(w_i^{(t)})$
- Intrusion Detection Trigger:

$$\delta_i(t) = \begin{cases} 1, & \text{if anomaly detected} \\ 0, & \text{otherwise} \end{cases}$$

5. Optimization Objectives

The global optimization goals of the FDT framework are to minimize evacuation delay and maximize overall safety:

$$\min_{R_i(t)} \sum_{i=1}^N T_i^{\text{evac}}(t), \quad \max_{R_i(t)} \sum_{i=1}^N \text{Evacuation_Success}(i, t)$$

Justification: The mathematical formulation of the Federated Digital Twin (FDT) framework is designed to model and support a scalable, privacy-preserving, and intelligent emergency response system across distributed cultural heritage environments. At the core of this framework, each museum is equipped with a Local Digital Twin (LDT) that ingests sensor data ($S_i(t)$) and performs localized analytics ($A_i(t)$) to produce evacuation recommendations ($R_i(t)$) through a context-aware decision function f_i . This setup ensures that each institution can make rapid and autonomous safety decisions during emergencies based on real-time inputs such as visitor density, hazard proximity, and structural conditions. From a system-wide coordination perspective, Federated Learning plays a critical role by enabling local AI models to be trained on-site using operational data, without transmitting sensitive information to a central server. This approach allows the system to scale across multiple institutions while preserving institutional autonomy and ensuring legal compliance with data privacy standards. The aggregated global model w_{global} , computed through federated

averaging, captures shared knowledge across all museums, including hazard trends, evacuation patterns, and visitor behaviors under stress. This aggregated insight helps inform adaptive safety strategies and enhances system-wide preparedness.

The emergency response capability is mathematically constrained by the concept of evacuation success, where an evacuation is considered effective if completed within a safety threshold T_{safe} . This binary condition allows the system to quantify emergency response effectiveness and drive optimization objectives minimizing evacuation time and maximizing overall success rates across all institutions. Such metrics are vital for evaluating performance under varying conditions, such as congestion, fire propagation, or communication delays.

Furthermore, the security and privacy layer ensures the trustworthiness and robustness of the system. Mechanisms like differential privacy, encryption, and intrusion detection safeguard data integrity and protect critical infrastructure against cyber threats that may arise during crises. These features are essential in cultural heritage contexts where physical and digital assets are both highly valuable and vulnerable.

Overall, this mathematical validation demonstrates that the FDT framework is capable of supporting real-time, decentralized decision-making, coordinated learning, and secure data handling across multi-site museum networks. It provides the computational backbone to guide adaptive evacuation protocols, enhance situational awareness, and safeguard human life and heritage artifacts during emergencies.

Acknowledgements

I would like to express my sincere gratitude to Professor Henry Muccini for his continuous guidance, scientific rigor, and invaluable support throughout the development of this doctoral research. His mentorship has been essential in shaping both the methodological foundations and the scientific contributions of this dissertation.

I would also like to thank the entire FrameLab team for their collaborative spirit, interdisciplinary expertise, and constructive discussions, which greatly enriched the scientific depth and practical relevance of this work. Their dedication to advancing Digital Twins, Artificial Intelligence, and Extended Reality for cultural heritage has been a constant source of motivation.

My appreciation also extends to the University of L'Aquila, whose academic environment and institutional support have enabled the successful advancement of this work. I gratefully acknowledge the contribution of the Italian National Recovery and Resilience Plan (PNRR), whose strategic commitment to digital innovation and cultural heritage preservation created the framework within which this research was conducted.

Finally, my sincere appreciation goes to the broader university community for fostering a supportive, stimulating, and intellectually vibrant environment that encouraged continuous learning and scientific growth throughout my doctoral journey.

Tesi redatta con il cofinanziamento dell'Unione europea-*NextGeneration EU*
Piano Nazionale di Ripresa e Resilienza Missione 4 – Componente 1 – Investimento 4.1 “Dottorati per la
Pubblica Amministrazione” – CUP E11I22000130001



Finanziato
dall'Unione europea
NextGenerationEU



Ministero
dell'Università
e della Ricerca



Italiadomani
PIANO NAZIONALE
DI RIPRESA E RESILIENZA



UNIVERSITÀ
DEGLI STUDI
DELL'AQUILA