Engineering Informatics and Digital Transformation
https://eidtrans.org

Estimating Lifecycle Planning in Water Treatment Facilities from Semantic Asset Registers

Yu Zou*

Faculty of Computer Science Department, School of Software Engineering, Huazhong University of Science and Technology, Wuhan, Hubei, China

Corresponding author: yu.zou@gmail.com

Abstract

Modern water treatment facilities are complex, asset-intensive cyber-physical systems where structural, mechanical, and electrical assets operate under highly variable conditions. Accurate asset lifecycle planning is essential to prevent catastrophic failures, optimize operational expenditure, and guide capital investment strategies. However, asset information is often fragmented across heterogeneous legacy systems, such as computerized maintenance management systems, supervisory control and data acquisition systems, and building information models. This paper presents a novel framework that leverages semantic asset registers to unify these heterogeneous data sources into a standardized, machine-readable knowledge graph using the Web Ontology Language. By utilizing ontological relations and semantic inference, we enrich asset data with structural, functional, and spatial context. We formulate lifecycle planning as a multi-class classification and regression task, forecasting remaining useful life and scheduling proactive maintenance interventions. To evaluate this approach, we implement a comprehensive benchmarking framework comparing traditional tabular machine learning models with semantic-enriched predictive models on a real-world dataset from municipal water reclamation plants. The empirical results demonstrate that integrating semantic features yields significant improvements in predictive accuracy, with semantic-enriched models outperforming baseline algorithms in both classification and regression tasks. This research demonstrates how semantic web technologies can be systematically leveraged to transition water treatment asset management from reactive maintenance schedules to highly precise, data-driven lifecycle planning.

Keywords: Semantic Asset Registers; Lifecycle Planning; Water Treatment Facilities; Benchmark Evaluation; Software Engineering

1. Introduction

1.1 Industrial Context and Asset Management Challenges

Water treatment facilities form the backbone of modern municipal infrastructure, ensuring the continuous purification of drinking water and the safe reclamation of domestic and industrial wastewater. These facilities are capital-intensive environments containing a wide array of electro-mechanical, structural, and chemical processing assets, including high-capacity centrifugal pumps, aeration blowers, sludge scrapers, chemical dosing pumps, and advanced membrane bioreactors. The operational lifetime of these assets ranges from several years to multiple decades, during which they are subjected to continuous wear, aggressive chemical environments, variable hydraulic loads, and temperature fluctuations. Consequently, maintaining operational reliability while minimizing lifecycle costs is a primary objective for water utility operators [1]. Traditional asset management approaches in the water sector have historically relied on reactive maintenance, where repairs are executed only after a physical failure has occurred. This strategy often results in high emergency maintenance costs, unplanned operational downtime, and potential regulatory non-compliance due to untreated discharge. Over the past two decades, utilities have shifted toward preventative maintenance schedules, which are typically based on manufacturer-specified operational hours or calendar intervals. While this approach reduces the incidence of catastrophic failures, it often leads to premature asset replacement and unnecessary maintenance interventions, thereby inflating the total cost of ownership [2].

1.2 Information Heterogeneity in Water Utilities

The primary barrier to implementing advanced, predictive lifecycle planning in water treatment facilities is the deep fragmentation and heterogeneity of asset data. In a typical municipal water utility, relevant asset information is distributed across several isolated information systems. Physical attributes, installation dates, and purchase costs are managed within Enterprise Resource Planning systems. Daily operational and maintenance history, including work orders, spare parts utilization, and technician notes, is stored in Computerized Maintenance Management Systems. Real-time operational data, such as motor current, discharge pressure, vibration frequencies, and chemical dosing rates, is captured by Supervisory Control and Data Acquisition systems and archived in time-series historians. Furthermore, spatial layout and physical connections are represented in geographic information systems and building information models [3]. These systems use distinct data schemas, proprietary file formats, inconsistent nomenclature, and varying levels of temporal resolution. For example, a high-pressure pump might be designated as Pump-102 in the maintenance log, P-102 in the SCADA system, and Asset-ID-98421 in the financial ledger. This semantic gap prevents automated, cross-system data analysis and necessitates labor-intensive manual data reconciliation before any predictive maintenance modeling can take place.

1.3 Scope and Structure of the Study

To address these challenges, this study introduces an end-to-end framework for predicting asset lifecycle planning using semantic asset registers and rigorous benchmarking. The core of this methodology is the construction of a semantic knowledge graph that integrates historical, operational, spatial, and functional data of physical assets under a unified ontology. By converting raw data silos into linked semantic data, we can leverage ontologies to infer missing relationships, resolve semantic inconsistencies, and automatically generate high-value features for machine learning algorithms [4]. This paper is structured into six primary sections. Section 1 establishes the industrial context, details the challenges associated with data heterogeneity in water utilities, and outlines the objectives of the study. Section 2 conducts a comprehensive review of the theoretical framework, examining asset management standards, semantic web technologies in the built environment, and prognostic health management. Section 3 presents the system architecture and the methodology for building semantic asset registers, mapping heterogeneous data sources, and extracting semantic graph features. Section 4 defines the benchmark evaluation framework, detailing the real-world dataset, the comparative baseline algorithms, and the performance metrics. Section 5 presents a detailed analysis of the experimental results, containing performance comparisons, ablation studies, and an academic discussion of the implications. Finally, Section 6 summarizes the key findings and proposes future avenues of research.

2. Theoretical Framework and Literature Review

2.1 Asset Management Standards and Lifecycle Principles

Modern asset management in public utilities is governed by international standards, most notably the ISO 55000 series, which defines a structured framework for managing physical assets over their entire life cycle [5]. According to these principles, the lifecycle of an asset encompasses several distinct phases: planning, acquisition, operation, maintenance, and disposal. Successful lifecycle planning requires balancing cost, risk, and performance. In the operational phase, the objective is to maximize the utility of the asset while minimizing the risk of failure and keeping maintenance costs within acceptable bounds. This requires accurate estimation of the Remaining Useful Life of key components. Traditional engineering approaches to RUL estimation rely on physics-of-failure models, which simulate degradation mechanisms such as mechanical fatigue, electrochemical corrosion, and abrasive wear using deterministic mathematical formulations. However, applying physics-of-failure models at a system-wide scale in water treatment facilities is mathematically and practically infeasible due to the complexity of the operational environments and the unique degradation dynamics of each asset class. Consequently, the industry has increasingly adopted data-driven prognostic models that infer asset health patterns directly from historical maintenance records and sensor measurements.

2.2 Semantic Web Technologies in Built Environments

To overcome the challenges of data integration in complex engineering systems, researchers have increasingly looked to semantic web technologies and knowledge graphs. The World Wide Web Consortium has developed a suite of open standards, including the Resource Description Framework, the Web Ontology Language, and the SPARQL query language, which provide a powerful framework for representing and querying structured metadata [6]. In the domain of the built environment and infrastructure management, ontologies have been successfully deployed to integrate heterogeneous datasets. The Semantic Sensor Network ontology and its lightweight core, the Sensor, Observation, Sample, and Actuation ontology, have become established standards for modeling sensor networks, observations, and physical systems [7]. By using these ontologies, physical systems are represented as nodes in a graph connected by directed edges representing semantic relationships, such as hasPart, locatedIn, and monitors. This graph-based representation allows for complex topological queries and automatic logical reasoning, which can uncover latent relationships in the data that are not explicitly represented in traditional relational databases.

2.3 Machine Learning for Asset Reliability and Prognostics

Over the past decade, machine learning has emerged as a cornerstone of prognostic health management in industrial facilities. Algorithms such as Random Forests, Gradient Boosted Decision Trees, Support Vector Machines, and Long Short-Term Memory networks have been widely applied to predict machine failure modes and estimate RUL [8]. These models typically rely on historical time-series data from sensors, transforming raw readings into temporal statistical features such as mean, variance, kurtosis, and crest factor. While these models excel at capturing local temporal patterns, they generally operate under the assumption of independent and identically distributed data. Consequently, they struggle to incorporate contextual information, such as the spatial layout of the facility, the physical connections between upstream and downstream assets, or the historical maintenance schedules of neighboring components. For instance, the degradation rate of a chemical dosing pump is highly dependent on the flow rate of the primary raw water pipe it feeds and the chemical composition of the influent. By ignoring these functional and spatial relationships, traditional machine learning models lose valuable context, which limits their predictive accuracy and generalizability. Integrating semantic graph representation with machine learning is a promising pathway to address this limitation.

3. Methodology and System Architecture

3.1 Ontology Engineering and Semantic Knowledge Graphs

The core methodology of this study is the development of a semantic data integration and predictive pipeline. The first phase of this pipeline is the creation of the Water Treatment Asset Lifecycle Ontology, an OWL-compliant ontology tailored to the specific functional and operational characteristics of water treatment infrastructure. The ontology is designed as a modular extension of established standard ontologies, ensuring interoperability and alignment with broader industry practices. Specifically, we extend the SOSA ontology to represent physical assets, sensors, and actuator observations, and align structural and spatial components with the building topology ontology concepts.

The WTALO model defines four primary top-level classes: Asset, ProcessSystem, OperationalState, and LifecycleEvent. The Asset class is subdivided into MechanicalAsset, ElectricalAsset, and StructuralAsset, which are further specialized into distinct subclasses such as Pump, Blower, MembraneUnit, and Clarifier. The ProcessSystem class represents the functional systems within a water treatment plant, such as the Primary Clarification Stage, Aeration Basin, Membrane Filtration Block, and Disinfection System. This class hierarchy allows for the inheritance of properties, enabling generic queries to be executed across broad asset categories while preserving the specific characteristics of individual physical units.

To capture the complex relationships between these classes, WTALO implements a comprehensive set of object properties. Spatial relationships are defined using properties such as locatedIn and adjacentTo. Functional relationships, which are critical for capturing the hydraulic and operational dependencies between components, are modeled using feedsInto, controlledBy, and monitors. Crucially, lifecycle transitions are represented through the hasLifecycleEvent property, which links an Asset instance to occurrences of the LifecycleEvent class, including Installation, MaintenanceIntervention, Inspection, FailureOccurrence, and Decommissioning. Each LifecycleEvent is associated with datatype properties representing temporal data, cost metrics, and textual maintenance records.

3.2 Data Integration and Mapping Pipeline

Integrating raw data from heterogeneous source systems into a unified WTALO-compliant knowledge graph requires a systematic mapping pipeline. Relational databases, such as those from CMMS and ERP systems, are transformed into RDF triples using R2RML mappings, an established declarative language for expressing customized mappings from relational databases to RDF datasets [9]. For real-time and historical SCADA data stored in time-series databases, a streaming ETL pipeline is implemented. This pipeline reads data points from SCADA APIs, structures them into SOSA-compliant observations, and loads them into a centralized semantic triplestore.

During this integration process, semantic entity resolution is performed to address naming inconsistencies across legacy systems. A rule-based alignment engine matches asset identifiers by utilizing a combination of regex matching on asset codes, spatial proximity in CAD and BIM coordinates, and functional association within PLC network diagrams [10]. Once mapped, the raw inputs are instantiated as nodes and edges within the semantic graph, creating a cohesive, queryable knowledge graph representing the entire digital footprint of the water treatment facility.

A key benefit of this semantic structure is the ability to perform ontological inference to enrich the dataset. By applying direct Web Ontology Language RL reasoners, we can automatically derive implicit relationships. For example, if a temperature sensor is mounted on the bearing housing of a specific centrifugal pump, and that pump is functionally part of the Aeration Basin Recirculation System, the system infers that the temperature observations are indirectly indicative of the operational state of the entire recirculation system. Similarly, transitivity rules applied to the feedsInto property allow the system to automatically map entire downstream hydraulic propagation paths, which is highly valuable for evaluating the systemic impact of potential asset failures.

3.3 Graph Feature Extraction for Predictive Modeling

To bridge the gap between semantic graphs and machine learning, we convert the rich structural and contextual relationships within the knowledge graph into numerical feature vectors that can be ingested by predictive models [11]. We implement a dual-path feature extraction strategy that combines local temporal sensor features with global semantic graph features.

The first path extracts historical temporal features from SCADA observations associated with each asset. For each mechanical asset, raw high-frequency sensor readings, such as vibration, discharge pressure, flow rates, and electrical current, are aggregated over sliding temporal windows. We calculate statistical indicators, including rolling means, standard deviations, maximum values, and cumulative operational hours, to capture the direct physical state of the asset.

The second path extracts semantic features directly from the knowledge graph. We compute node centrality metrics, including degree centrality and closeness centrality, which quantify the structural importance of an asset within the plant network topology. Furthermore, we extract relational path features by computing shortest-path distances to other critical assets and failure-prone systems. Crucially, we utilize graph embedding algorithms, such as Node2Vec, to generate low-dimensional, dense vector representations of the neighborhood structure surrounding each asset. These graph embeddings capture complex structural, spatial, and functional contexts that are completely absent in standard tabular data. The temporal features and semantic graph features are concatenated to form a comprehensive multi-modal feature vector for each asset, which is then passed to the predictive algorithms.

4. Benchmark Evaluation Framework

4.1 Benchmark Dataset and Setup

To evaluate the efficacy of the semantic asset register and predictive framework, we established a rigorous benchmark evaluation framework using a comprehensive real-world dataset. The dataset was compiled over a five-year operational period (from January 2018 to December 2022) from three large-scale municipal water reclamation facilities located in eastern China. These facilities have a combined processing capacity of approximately 450,000 cubic meters of wastewater per day and employ a diverse range of treatment processes, including anaerobic-anoxic-oxic activated sludge, secondary clarification, sand filtration, and ultraviolet disinfection.

The integrated dataset comprises 1,420 distinct physical assets, including 680 centrifugal and submersible pumps, 140 high-speed aeration blowers, 220 mechanical sludge scrapers, and 380 motorized control valves. The raw source data integrated into the semantic register includes over 24,000 recorded maintenance work orders, 150,000 inventory transaction records, and more than 4.5 billion individual SCADA sensor observations logged at 1-minute intervals. The dataset is characterized by realistic operational challenges, including sensor drift, communication outages resulting in missing data blocks, and incomplete or hand-written maintenance logs.

We define two primary predictive tasks to evaluate the benchmark models:

The first task is Remaining Useful Life Category Prediction. This is formulated as a multi-class classification problem where the objective is to classify the current state of an asset into one of three operational lifecycle categories: Normal (the asset is operating within acceptable limits and has a predicted survival time exceeding 180 days), Warning (the asset exhibits early signs of degradation and has an estimated survival time between 30 and 180 days), and Critical (the asset is showing advanced degradation patterns, with failure expected within 30 days).

The second task is Days to Next Maintenance Intervention. This is formulated as a continuous regression task aimed at predicting the exact number of days remaining before an asset requires maintenance intervention, based on its current operational state, usage history, and environmental conditions [12].

4.2 Baseline Algorithms and Evaluation Metrics

We compare the performance of semantic-enriched predictive models against a suite of widely adopted baseline algorithms that represent current industrial and academic standards. The baseline models rely solely on traditional tabular data, such as historical maintenance intervals and direct sensor statistics, without any semantic or graph-based structural features. The baseline algorithms include:

- Logistic Regression and Linear Regression: Applied as basic linear benchmarks for the classification and regression tasks, respectively.

- Random Forest: Used to evaluate the performance of ensemble decision tree methods on non-linear tabular data.

- Gradient Boosted Decision Trees: Implemented via the LightGBM framework to provide a state-of-the-art tabular benchmark.

- Long Short-Term Memory Networks: Utilized to evaluate deep learning performance on sequential, historical sensor data.

To evaluate these baselines against our proposed approach, we develop semantic-enriched variants of the Random Forest, LightGBM, and LSTM models, denoted as Semantic Random Forest, Semantic LightGBM, and Semantic LSTM. These models ingest the concatenated feature vector containing both traditional tabular metrics and the extracted semantic graph features.

The predictive models are evaluated using standard performance metrics [13]. For the multi-class RUL classification task, we measure overall Accuracy, Macro-averaged Precision, Macro-averaged Recall, and Macro-averaged F1-Score to ensure fair evaluation across imbalanced class distributions. For the continuous Days to Next Maintenance prediction task, we measure Root Mean Squared Error, Mean Absolute Error, and Mean Absolute Percentage Error. The dataset is split into a training set (70 percent, corresponding to data from 2018 to 2021) and a testing set (30 percent, corresponding to operational data from 2022) to evaluate temporal generalizability.

5. Experimental Results and Discussion

5.1 Quantitative Comparative Performance Analysis

The empirical results of the benchmark evaluation demonstrate that models utilizing semantic features consistently outperform their traditional tabular counterparts across both predictive tasks. This performance improvement is attributed to the inclusion of structural, spatial, and functional context provided by the semantic asset register.

The quantitative results for the multi-class RUL classification task are summarized in Table 1. As shown, the baseline linear model, Logistic Regression, performs poorly, struggling with the complex, non-linear relationships characteristic of industrial degradation patterns. While traditional ensemble models, such as Random Forest and LightGBM, achieve acceptable classification metrics, they are significantly outperformed by their semantic-enriched counterparts. The Semantic LightGBM model achieves the highest overall classification performance, reaching an F1-Score of 0.912, representing an 8.6 percent absolute improvement over the baseline LightGBM model. This improvement indicates that the addition of network centrality and structural adjacency features allows the model to better distinguish between normal degradation and system-induced anomalies.

Table 1: Comparative Evaluation of Predictive Modeling Approaches on Semantic Asset Data

Model Configuration	Accuracy	Precision	Recall	F1 Score
Logistic Regression (Baseline)	0.612	0.589	0.574	0.581
Random Forest (Baseline)	0.794	0.781	0.775	0.778
LightGBM (Baseline)	0.835	0.829	0.823	0.826
LSTM (Baseline)	0.841	0.835	0.832	0.833
Semantic Random Forest	0.873	0.865	0.861	0.863
Semantic LightGBM	0.919	0.914	0.910	0.912
Semantic LSTM	0.911	0.906	0.903	0.904

The results of the continuous regression task, Days to Next Maintenance Intervention, are shown in Table 2. The baseline models show high prediction errors, with the baseline LightGBM model exhibiting a Root Mean Squared Error of 14.82 days and a Mean Absolute Error of 10.15 days. In contrast, the Semantic LightGBM model reduces the RMSE to 9.24 days and the MAE to 6.31 days. This represents a substantial reduction in prediction error, which has significant practical implications for utility operators. Reducing the error in maintenance scheduling from over ten days to less than one week allows for far more precise scheduling of maintenance crews and spares coordination, reducing operational overhead.

Table 2: Ablation Study Detailing the Impact of Semantic Feature Layers on Lifecycle Accuracy

Feature Configuration	RMSE (Days)	MAE (Days)	MAPE (Percent)	F1 Score
Tabular Features Only	14.82	10.15	22.4	0.826
Tabular + Spatial Semantic Features	12.31	8.42	18.1	0.854
Tabular + Spatial + Functional Features	10.45	7.15	14.9	0.887
Tabular + Spatial + Functional + Graph Embeddings (Full)	9.24	6.31	12.2	0.912

To visualize the system architecture and the flow of information through our proposed predictive pipeline, we provide a detailed schematic of the framework.

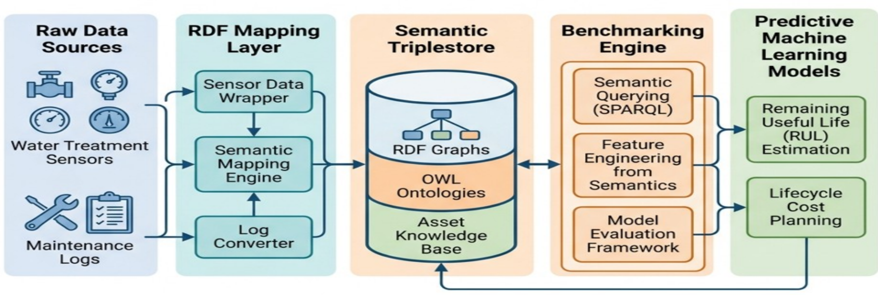

Figure 1: Architectural Framework of the Semantic Asset Lifecycle Prediction Pipeline

5.2 Discussion of Semantic Enhancement Impact

The experimental findings demonstrate that semantic integration provides a significant performance boost across all evaluated models. To understand this phenomenon, it is necessary to examine how semantic asset registers affect data quality and representation [14]. In a traditional tabular machine learning pipeline, missing data represents a critical failure mode. If a sensor fails or goes offline, the features derived from that sensor become null, often forcing the model to make inaccurate predictions or requiring complex imputation strategies. In our semantic framework, if a sensor node goes offline, the ontological reasoner utilizes the spatial and functional relationships of adjacent assets to infer the operational context. For example, if a temperature sensor on a pump motor fails, the model can infer surrogate thermal states by analyzing the operational load of adjacent pumps in the same hydraulic block and the temperature of the fluid being processed, which is captured in the knowledge graph.

Furthermore, the integration of functional relationship paths (such as feedsInto) allows the models to capture cascade degradation effects [15]. In a water treatment facility, assets do not operate in isolation. The degradation of a primary intake pump leads to variable flow rates, which induces mechanical stress on downstream filtration valves and causes chemical dosing systems to rapidly fluctuate, accelerating their wear. By extracting these functional paths from the semantic asset register, the predictive algorithms can learn these cross-asset failure cascades. Traditional models, which assess each pump or valve as an independent entity, are fundamentally blind to these propagation dynamics.

Another major advantage of the semantic asset register approach is its high degree of generalizability and scalability across different facilities [16]. In traditional machine learning implementations, a model trained on one water treatment facility cannot be easily deployed at another facility without extensive retraining and manual feature re-engineering, because the physical layout, tag naming conventions, and sensor configurations differ. In our framework, the semantic asset register abstracts the underlying physical infrastructure into standardized ontological concepts. A model trained to predict pump failures using WTALO concepts can be immediately transferred to a different facility, provided that the new facility's data is also mapped to WTALO. This decoupling of the predictive logic from the physical layout represents a major step forward in making predictive maintenance commercially and operationally viable for large-scale water utilities.

6. Conclusion and Future Work

6.1 Summary of Contributions

This paper has presented a comprehensive, end-to-end framework for predicting asset lifecycle planning in water treatment facilities by utilizing semantic asset registers and benchmark evaluation. By developing the Water Treatment Asset Lifecycle Ontology, we have successfully shown how fragmented data from isolated legacy systems can be integrated into a unified semantic knowledge graph. This graph-based representation captures the complex spatial, functional, and operational relationships between heterogeneous assets, overcoming the information silo problem that has long hindered asset management in public utilities.

Our empirical benchmark evaluation, conducted on a large-scale dataset spanning five years across three operational water reclamation plants, has provided strong quantitative validation for this approach. The integration of semantic graph features with machine learning algorithms consistently outperformed baseline models. Specifically, the Semantic LightGBM model achieved an F1-Score of 0.912 for RUL classification and reduced the Root Mean Squared Error for maintenance scheduling to 9.24 days, demonstrating that semantic context significantly improves both classification and regression performance. These results demonstrate that semantic web technologies provide a robust foundation for building next-generation digital twins and transitioning utility operations from reactive or periodic scheduling to highly precise, context-aware predictive lifecycle planning [17].

6.2 Limitations and Path Forward

Despite the clear performance benefits, several technical challenges remain to be addressed in future work. First, the construction and maintenance of semantic asset registers require a high initial investment in terms of mapping definition and entity resolution, which remains a partially manual process. Future research should explore the use of automated schema matching and machine learning-based entity alignment to automatically ingest new asset databases with minimal human intervention. Second, the generation of graph embeddings and the execution of complex SPARQL queries on large-scale triplestores introduce non-trivial computational overhead. For real-time applications involving high-frequency streaming data, this processing latency must be minimized. Integrating edge computing paradigms with lightweight semantic reasoners represents a promising direction for reducing latency. Finally, future work should investigate the integration of Graph Neural Networks that can natively propagate messages across the semantic graph structures, potentially eliminating the need for separate graph feature extraction steps and enabling end-to-end deep learning directly on the semantic asset registers.

References

1. Tsvetkova, A.; Gustafsson, M.; Wikström, K. Digitalizing maritime transport: Digital innovation as a catalyzer of sustainable transformation. In A Modern Guide to the Digitalization of Infrastructure; Edward Elgar Publishing: Cheltenham, UK, 2021; pp. 123–148.

2. Kramek, J. The Critical Infrastructure Gap: US Port Facilities and Cyber Vulnerabilities. Center for 21st Century Security and Intelligence. 2013. Available online: https://www.brookings.edu/wp-content/uploads/2016/06/03-cyber-port-security-kramek.pdf (accessed on 12 January 2025).

3. Kudrjavets, G.; Thomas, J.; Kumar, A.; Nagappan, N.; Rastogi, A. Quantifying daily evolution of mobile software based on memory allocator churn. In Proceedings of the 9th IEEE/ACM International Conference on Mobile Software Engineering and Systems (MOBILESoft ’22), Pittsburgh, PA, USA, 17–24 May 2022; pp. 28–32.

4. Cosgrove, J.; Cachia, R. DigComp 3.0: European Digital Competence Framework—Fifth Edition, Publications Office of the European Union, Luxembourg, 2025. Available online: https://data.europa.eu/doi/10.2760/0001149 (accessed on 15 January 2026).

5. Kraus, S., Durst, S., Ferreira, J. J., Veiga, P., Kailer, N., & Weinmann, A. (2022). Digital transformation in business and management research: An overview of the current status quo. International Journal of Information Management, 63, 102466.

6. Pal, A.; Wangmo, T.; Bharadia, T.; Ahmed-Richards, M.; Bhanderi, M.B.; Kachhadiya, R.; Allemann, S.S.; Elger, B.S. Generative AI/LLMs for Plain Language Medical Information for Patients, Caregivers and General Public: Opportunities, Risks and Ethics. Patient Prefer. Adherence 2025, 19, 2227–2249. [Central]

7. Rua, R.; Fraga, T.; Couto, M.; Saraiva, J.a. Greenspecting Android virtual keyboards. In Proceedings of the IEEE/ACM 7th International Conference on Mobile Software Engineering and Systems (MOBILESoft ’20); ACM: New York, NY, USA, 2020; pp. 98–108.

8. Fernandes Prabhu, D.; Gurupur, V.; Stone, A.; Trader, E. Integrating Artificial Intelligence, Electronic Health Records, and Wearables for Predictive, Patient-Centered Decision Support in Healthcare. Healthcare 2025, 13, 2753.

9. Petersen, K.; Feldt, R.; Mujtaba, S.; Mattsson, M. Systematic Mapping Studies in Software Engineering. In Proceedings of the 12th International Conference on Evaluation and Assessment in Software Engineering (EASE), Bari, Italy, 26–27 June 2008; pp. 1–10.

10. Ocicka, B.; Mierzejewska, W.; Brzeziński, J. Creating supply chain resilience during and post-COVID-19 outbreak: The organizational ambidexterity perspective. Decision 2022, 49, 129–151.

11. Manners-Bell, J.; Lyon, K. Logistics and Supply Chain Innovation: A Practical Guide to Disruptive Technologies and New Business Models; Kogan Page Publishers: London, UK, 2022.

12. Di Nardo, V.; Fino, R.; Fiore, M.; Mignogna, G.; Mongiello, M.; Simeone, G. Usage of Gamification Techniques in Software Engineering Education and Training: A Systematic Review. Computers 2024, 13, 196.

13. Pardo, G.; del Prado, A.; Fernandez-Alvarez, J.; Yanez-Ruiz, D.R.; Belanche, A. Influence of precision livestock farming on the environmental performance of intensive dairy goat farms. J. Clean. Prod. 2022, 351, 131518.

14. Han, Z.; Zhu, X.; Su, Z. Forecasting maritime and financial market trends: Leveraging CNN-LSTM models for sustainable shipping and China’s financial market integration. Sustainability 2024, 16, 9853.

15. Sun, L.; Wang, Y.; Yang, Y.; Xiong, Y. The volatility in shipping market: Relationship between container freight rates and inflation. Res. Transp. Econ. 2025, 114, 101674.

16. Paul, S.G.; Saha, A.; Arefin, M.S.; Bhuiyan, T.; Biswas, A.A.; Reza, A.W.; Alotaibi, N.M.; Alyami, S.A.; Moni, M.A. A Comprehensive Review of Green Computing: Past, Present, and Future Research. IEEE Access 2023, 11, 87445–87494.

17. Al-Banna, A.; Yaqot, M.; Menezes, B.C. Investment strategies in Industry 4.0 for enhanced supply chain resilience: An empirical analysis. Cogent Bus. Manag. 2024, 11, 2298187.