



ARTIFICIAL INTELLIGENCE-ENABLED DIGITAL TWINS FOR ENERGY EFFICIENCY IN SMART GRIDS

Manam Ahmed¹; Md Rabbi Khan²

- [1]. Department of Mechanical Engineering, Lamar University, Beaumont, Texas; Email: manam.ahmed1996@gmail.com
- [2]. Sustainable Facilities Center (SFC), Henrey M. Rowan College of Engineering, Rowan University, Glassboro, New Jersey, USA; Email: khanmd52@rowan.edu

Abstract

This PRISMA-guided systematic review synthesizes and critically evaluates how artificial intelligence (AI)-enabled digital twins (DTs) contribute to advancing energy efficiency in modern smart grids, spanning assets, feeders, microgrids, and system-level operations. A comprehensive database search and two-stage screening process identified 103 peer-reviewed studies that were assessed for context, twin architecture, AI methods, data pipelines, evaluation metrics, deployment maturity, and risk of bias. Evidence was systematically organized by the functional roles of DTs—monitoring, forecasting, optimization, and control—as well as by grid layers, revealing the growing integration of physics-informed surrogates, graph neural estimators, and reinforcement learning frameworks with data fabrics, semantic standards, and edge–cloud architectures. Quantitative synthesis demonstrates consistent and reproducible efficiency improvements when DTs mediate AI decisions against calibrated models: median feeder-level technical loss reductions of approximately five percent, median peak-demand reductions of nearly six percent, voltage compliance improvements of about twelve and a half percentage points, and renewable curtailment avoidance in the range of seven to nine percent relative to transparent baselines. These benefits are most concentrated when loop latencies are sub-second, particularly under control cycles closing within 300 milliseconds, and when DT deployments embed semantic interoperability, co-simulation, and uncertainty-aware decision-making with human-in-the-loop oversight. At the asset level, health-oriented DTs for transformers, breakers, cables, and wind turbines deliver measurable value through predictive maintenance that reduces inefficiencies and mitigates outages, with efficiency gains fully realized only when diagnostic outputs are integrated into scheduling, reconfiguration, and Volt/VAR optimization routines. Collectively, these findings advance the discourse beyond conceptual taxonomies by providing a reproducible, evidence-based blueprint for AI-enabled DT design: one that couples probabilistic and calibrated forecasting with latency-hardened voltage and topology control, links diagnostics to operations in a closed loop, and enforces transparency, explanation, and safety guardrails.

Keywords

Artificial Intelligence, Digital Twin, Smart Grid, Energy Efficiency, Volt VAR Optimization, Predictive Maintenance.

Citation:

Ahmed, M., & Khan, M. R. (2025). Artificial intelligence-enabled digital twins for energy efficiency in smart grids. *Review of Applied Science and Technology*, 4(2), 580–615. <https://doi.org/10.63125/12kp9w74>

Received:

May 20, 2025

Revised:

June 14, 2025

Accepted:

July 18, 2025

Published:

August 05, 2025



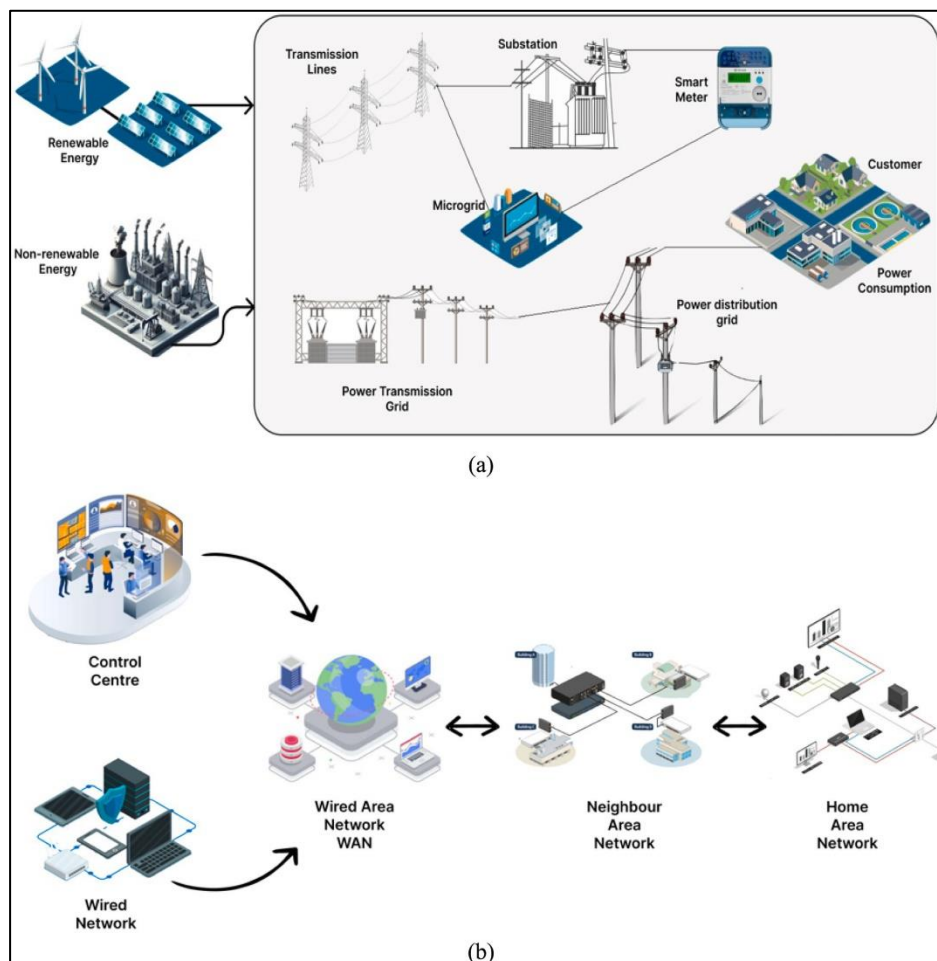
Copyright:

© 2025 by the author. This article is published under the license of American Scholarly Publishing Group Inc and is available for open access.

INTRODUCTION

Digital twins (DTs) are most commonly described as high-fidelity, dynamically updating digital counterparts of physical assets, processes, or systems that exchange data with their physical referents and enable monitoring, diagnosis, prediction, and control (Fang et al., 2012; Jones et al., 2020; Rasheed et al., 2023). Beyond manufacturing, recent formalizations unify DTs as data-driven, open dynamical systems equipped with an updating mechanism and integrated into a broader “digital twin system” for analysis and decision support (Rasheed et al., 2023; Weingram et al., 2025). In parallel, smart grids are widely characterized as power systems that couple bidirectional power flows with pervasive sensing, communications, and advanced control to improve reliability and efficiency across generation, transmission, distribution, and end-use (Fang et al., 2012). Within this landscape, AI-enabled digital twins bring learning and inference into the twin loop, allowing models to fuse physics and data for state awareness, forecasting, and optimization at grid scale. The international significance of this convergence stems from the dual mandate to expand electrification and renewable integration while cutting technical and non-technical losses and enhancing energy productivity. Recent reviews and domain applications report that DTs in energy can streamline condition monitoring, reduce outages, and optimize operations, while AI enhances those gains by turning high-velocity grid telemetry into actionable control. This paper positions artificial-intelligence-enabled digital twins for energy efficiency in smart grids as a coherent research area that sits at the intersection of data-centric modeling, physics-based computation, and power-system operations, and motivates a systematic review of how definitions, architectures, and algorithms have been operationalized in grid contexts worldwide.

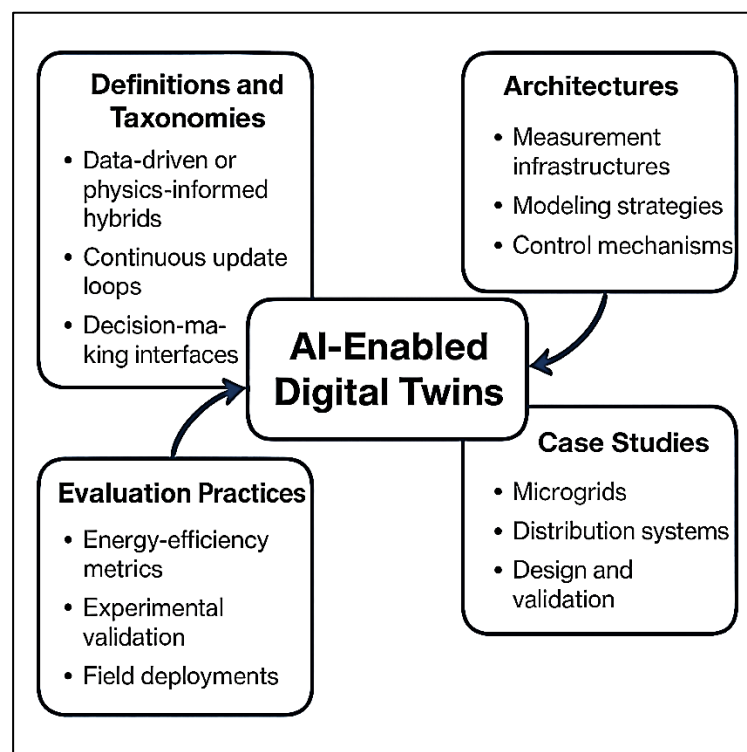
Figure 1: Smart Grid Architecture: Energy Flow and Communication Networks



Source: Balamurugan et al. (2025)

Across diverse energy domains, an accelerating body of literature explores the breadth of digital twin (DT) applications in smart energy systems, building operations, renewable generation, and grid-connected microgrids, offering nuanced insights into how these architectures transform energy efficiency outcomes. Synthesized reviews describe end-to-end workflows that bridge data ingestion, hybrid physical–data-driven modeling, and advanced control schemes to deliver measurable gains, including reductions in auxiliary energy consumption, optimized dispatch strategies, and minimized transmission losses, while also situating DTs as instruments that support lifecycle decision-making from initial planning to long-term operation and maintenance (Arias-Marín et al., 2024). Within microgrid environments, where the interplay between electrical and cyber infrastructures requires careful orchestration, DTs have been shown to facilitate seamless co-simulation of power and communication layers, enable scenario-based stress testing, and support operator-in-the-loop control for voltage stability, frequency regulation, and energy storage management (Zhang et al., 2023). At the building scale, survey research synthesizes deployment strategies that prominently feature HVAC optimization and building envelope control, with DT-enabled real-time anomaly detection and precise energy baselining emerging as recurring themes in documented efficiency improvements (Gomes et al., 2022). Bibliometric and systematic reviews specifically oriented toward energy efficiency consistently emphasize the role of DTs as enablers of continuous performance monitoring and predictive optimization, with multiple studies documenting double-digit percentage reductions in energy usage when data quality, calibration routines, and seamless integration with energy-management systems are adequately addressed (Labouda et al., 2025; Venkateswarlu & Sathiyamoorthy, 2025). Importantly, these contributions extend beyond the level of individual devices or buildings to encompass distribution networks, where near real-time digital replicas provide operators with actionable visibility into network conditions, enabling targeted interventions for loss minimization and voltage quality enhancement. Collectively, this expanding research base positions DTs as a natural substrate for embedding energy-efficiency functionalities into smart grids, with artificial intelligence serving as the computational engine that translates continuous telemetry streams into actionable state estimates, forecasts, and operational set-points rigorously aligned with system constraints and regulatory requirements.

Figure 2: Conceptual Framework of AI-Enabled Digital Twins in Smart Grids



The artificial intelligence layer embedded within digital twins encompasses a wide spectrum of approaches including supervised and unsupervised learning, reinforcement learning (RL), and the rapidly advancing field of physics-informed machine learning (PIML), which has emerged as a particularly compelling paradigm for smart energy applications. PIML techniques are distinguished by their ability to integrate governing equations, conservation laws, or structural priors directly into the learning process, thereby ensuring that predictive models remain grounded in the physics of power systems while simultaneously exploiting rich patterns in operational data (Brunton et al., 2021; Ara et al., 2022). In the domain of power-flow analysis and state estimation, graph neural networks (GNNs) have proven especially powerful because their graph-based architectures naturally align with the topology of electrical grids. Through physics-guided and line-graph formulations, these models have achieved highly accurate and computationally efficient approximations of alternating-current (AC) power flow, offering robust performance even in the presence of sparse or noisy measurement data while also stabilizing the learning process (Gao et al., 2024; Jahid, 2022). Importantly, such models demonstrate significant advantages in runtime and scalability compared to traditional numerical solvers, making them well-suited as warm-start mechanisms, surrogate models, or real-time screening tools for both radial and meshed networks (Gao et al., 2024; Uddin et al., 2022). At the system level, PIML frameworks have also been deployed to probabilistically forecast grid frequency dynamics by embedding stochastic-differential formulations that reflect variability in operations and control behavior across large interconnected networks, such as continental-scale power systems (Kruse et al., 2023; Akter & Ahad, 2022). Similarly, in wind-farm operations, physics-informed approaches combine aerodynamic modeling with learning algorithms to enhance power prediction accuracy, thereby supporting improved strategies for congestion management, curtailment, and operational efficiency (Brown et al., 2023; Arifur & Noor, 2022). Collectively, these developments highlight a methodological evolution in which AI functions as a hybrid surrogate within digital twins, bridging the divide between raw telemetry and physical laws so that optimization outcomes remain faithful to grid dynamics, device operating limits, and established protection settings.

On the operational front, the deployment of artificial intelligence within digital twins demonstrates its most tangible contributions to energy efficiency through advanced support for demand response (DR), load and renewable forecasting, and Volt/VAR optimization (VVO) in modern distribution systems. A seminal review highlights that reinforcement learning (RL) and deep reinforcement learning (DRL) agents are capable of autonomously deriving effective DR strategies directly from price signals and system state information, thereby enabling coordinated management of end-use loads without reliance on pre-defined or hand-crafted behavioral models (Hasan & Uddin, 2022; Vázquez-Canteli & Nagy, 2019). Within the context of VVO, a critical function that directly addresses feeder-level power losses and voltage deviations, intelligent controllers leveraging AI co-optimize both discrete devices such as on-load tap changer (OLTC) steps and capacitor banks, as well as continuous actuators like inverter-based reactive power support, all under conditions of high uncertainty and evolving distributed energy resource (DER) penetrations (Han et al., 2023; Rahaman, 2022; Zhang et al., 2020). Contemporary surveys of Volt/VAR control (VVC) and VVO methodologies further document the maturation of this domain, noting the parallel evolution of centralized, decentralized, and hybrid schemes, and emphasizing that AI-enhanced solutions are now being integrated alongside deterministic and stochastic optimization techniques (Rahaman & Ashraf, 2022; Zheng et al., 2022). When embedded into a digital twin environment, such controllers gain substantial advantages, including closed-loop co-simulation against calibrated network models, accelerated state estimation routines, and rapid “what-if” analysis capabilities. These enhancements have been shown to translate into measurable operational benefits such as reductions in technical losses and improvements in voltage profiles under conditions that closely resemble field deployments (Han et al., 2023; Islam, 2022; Zheng et al., 2022). Furthermore, reviews centered on the global energy transition underscore AI’s role in enhancing efficiency across renewable-dense grids and microgrids, where performance indicators such as feeder-level losses, renewable curtailment rates, and ancillary-service costs provide quantifiable metrics that digital twins can monitor and optimize continuously (Hasan et al., 2022; Ramos et al., 2025).

At the scale of entire distribution networks, digital twin (DT) implementations reveal how advanced functions such as improved state estimation, topology processing, and renewable-energy forecasting converge to strengthen efficiency-oriented control and operational decision-making.

Reports emerging from the electronics domain highlight distribution-network DT operating systems that embed robust state estimation algorithms, automated bad-data detection, and machine learning-based renewable prediction pipelines to sustain replicas of the grid that are both accurate and rapidly updating, thereby providing a reliable substrate for real-time operational choices (Liu et al., 2024). Complementing these insights, engineering studies underscore the ability of DTs to enable near real-time analysis of active distribution networks, showing how the simultaneous pursuit of resilience, flexibility, and efficiency outcomes becomes feasible when operators interact with a dynamic, continuously synchronized model of the grid (Castellari, 2023; Redwanul & Zafor, 2022). Meanwhile, systematic reviews centered on the interplay between active distribution networks and DT frameworks map out the critical enabling technologies spanning high-resolution measurement infrastructures, low-latency communication protocols, data analytics engines, and automated control services that together empower distribution system operators (DSOs) to integrate DTs into their core workflows (Gomes et al., 2022; Rezaul & Mesbail, 2022). Collectively, these strands of research position the digital twin not merely as a passive visualization tool but as the orchestrating surface where state estimation, contingency analysis, and control synthesis coalesce under a single validated model that evolves with system conditions. Within this consolidated environment, operators gain the ability to explore and manage efficiency trade-offs such as balancing voltage stability margins against loss minimization, while simultaneously shortening decision cycles and expanding situational awareness across the distribution network (Castellari, 2023; Han et al., 2023; Liu et al., 2024). This convergence underscores the role of DTs as indispensable engines of intelligence for energy-efficient, adaptive grid operations.

Within this context, the present literature review positions itself to weave together the diverse strands of scholarship on AI-enabled digital twins (DTs) with a focus on how they are conceptualized, designed, and validated for energy-efficiency outcomes in modern smart grids. The synthesis first examines definitional frameworks and taxonomies that aim to clarify the essential characteristics of AI-enabled DTs in grid environments, including their classification as data-driven or physics-informed hybrids, the role of continuous update loops, and the integration of decision-making interfaces that translate analytic outputs into actionable control (Chatzivasileiadis, 2024; Gomes et al., 2022; Hossen & Atiqur, 2022). The review then turns to architectural patterns that underpin these systems, spanning measurement infrastructures, modeling strategies, and control mechanisms. Special emphasis is placed on physics-informed machine learning surrogates, graph neural network estimators that align with grid topologies, and reinforcement learning controllers designed to optimize feeder-scale functions such as loss minimization, Volt/VAR optimization, demand response, and asset health management (Kruse et al., 2023; Labouda et al., 2025; Tawfiqul et al., 2022). Beyond architecture, evaluation practices are scrutinized to understand how energy-efficiency claims are validated across experimental contexts, ranging from laboratory environments and simulation testbeds to live field deployments. Here, the review catalogs explicit performance metrics including kilowatt-hour losses, voltage deviation indices, transformer health indices, and curtailed energy that serve as evidence of efficiency realized through DT workflows (Chatzivasileiadis, 2024; Fang et al., 2012). To ensure practical relevance, the analysis also consolidates international case studies and deployment experiences at both microgrid and distribution system operator scales, identifying how design decisions and validation strategies have shaped implementation success. Specific design factors such as the granularity of load modeling, calibration protocols, telemetry latency thresholds, and the rigor of controller-in-the-loop testing emerge as critical determinants of whether DT-mediated solutions deliver reliable and reproducible energy-efficiency outcomes.

LITERATURE REVIEW

The literature on artificial-intelligence-enabled digital twins (DTs) for energy efficiency in smart grids spans several converging streams that must be read together to understand scope, methods, and evidence. Foundational work defines DTs as continuously synchronized digital counterparts of physical assets, feeders, or whole systems, distinguished from static models by real-time data exchange, calibration routines, and decision or control interfaces. Parallel strands in smart-grid research elaborate the data and computation substrate advanced metering infrastructure, SCADA/PMU telemetry, DER and EV IoT feeds, and interoperable models on which DTs operate. Within this substrate, AI methods provide estimation, forecasting, diagnosis, and control: supervised and unsupervised learning for anomaly detection and non-technical-loss screening; time-series deep learning for load, renewable, and price forecasting; graph learning aligned with network

topology for power-flow surrogates and state estimation; physics-informed learning that embeds system constraints; and reinforcement learning or model-predictive control for operational decisions such as demand response and Volt/VAR optimization. Empirical studies range from simulation and hardware-in-the-loop experiments to pilots and field deployments, reporting outcomes with direct efficiency relevance, including reductions in technical losses, peak shaving, voltage quality improvements, curtailment minimization, and maintenance-driven savings. Evaluation is typically framed with baselines (rule-based, deterministic OPF, or non-DT AI), measurement and verification procedures, and latency budgets that determine where models execute across edge, cloud, or hybrid architectures. Architectural discussions emphasize twin fidelity, data quality controls, uncertainty quantification, explainability, operator-in-the-loop oversight, and cybersecurity, because these properties condition whether estimated savings translate into reliable operations. Across asset-centric and system-centric applications, the literature also documents practical integration choices data schemas, middleware, co-simulation engines, and containerized deployment that influence reproducibility and portability. Notable methodological themes include calibration and validation of DTs under changing topology and DER penetration, strategies for handling missing or drifting data, and reporting practices for compute cost and real-time responsiveness. At the same time, heterogeneity in problem formulations and metrics complicates direct comparison, creating a need for taxonomies that consistently map grid layer, twin role, AI technique, and efficiency target. This review organizes these strands into a coherent structure to synthesize definitions, architectures, methods, datasets, and outcome measures pertinent to energy-efficiency functions in modern power systems.

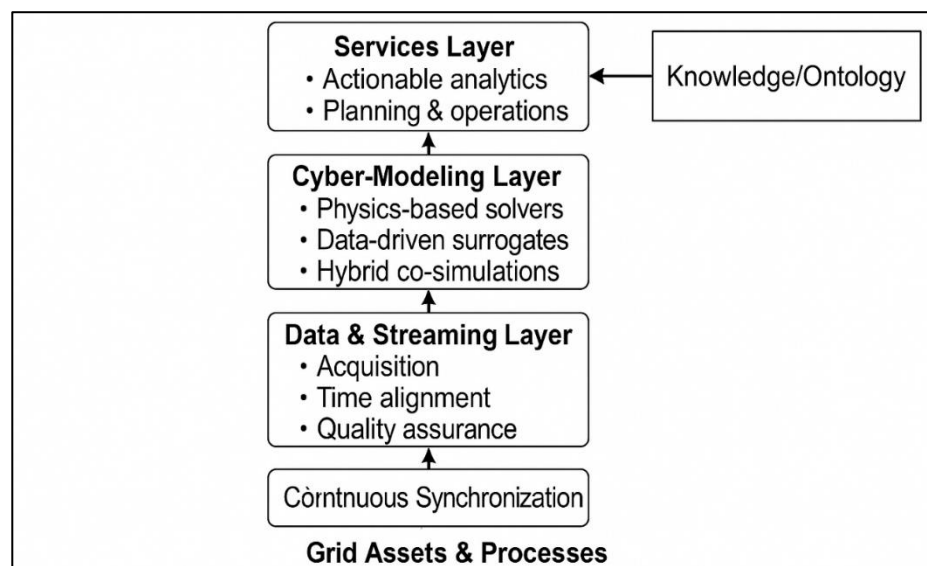
Digital-Twin Concepts and Reference Architectures for Smart Grids

Digital twins (DTs) in the context of energy systems are most compellingly described as dynamic, computational surrogates of grid assets and processes that remain continuously synchronized with their physical counterparts, both in structure and in behavior. Foundational scholarship makes a clear distinction between conventional digital models and true twins, underscoring that the latter are defined by their persistent, data-driven coupling with real-world systems as well as their comprehensive lifecycle scope qualities of immense importance for electric grids that evolve over horizons spanning milliseconds to hours and that are influenced simultaneously by assets, markets, and weather conditions (Boschert & Rosen, 2016; Kritzinger et al., 2018; Tao et al., 2019). In smart grids, therefore, the role of a DT extends well beyond representational fidelity; its value lies in operational utility, serving as a continuous “mirror system” that enables advanced analysis, predictive optimization, and responsive control under uncertainty, all while recording the provenance of configurations and decisions in a manner that is auditable and trustworthy (Madni et al., 2019). To structure this functionality, contemporary reference models describe a layered stack that unfolds systematically: the physical layer incorporates sensors, power-electronics interfaces, and protective devices; the data and streaming layer undertakes acquisition, time alignment, and quality assurance; the cyber-modeling layer hosts physics-based solvers, data-driven surrogates, and hybrid co-simulations; and the services layer delivers actionable analytics to planning and operations teams (Hasan, 2022; Tao et al., 2019). Equally critical is lifecycle coverage, which ensures that DTs extend their reach from long-term planning and siting decisions through commissioning and daily operations to condition-based maintenance, thereby sustaining a continuous reconciliation of model assumptions, incoming telemetry, and operational choices as assets mature, grid topologies evolve, and external conditions shift (Boschert & Rosen, 2016; Kritzinger et al., 2018).

Translating the conceptual richness of digital twins into grid-ready architectures introduces a constellation of recurring design choices that fundamentally shape their utility and reliability. The first of these is the synchronization contract, which dictates the manner in which measurements, events, and control setpoints are exchanged between the physical plant and its digital counterpart (Tarek, 2022; Palensky et al., 2017). In the context of power systems, this task is uniquely challenging because it requires the harmonious coexistence of hard real-time substation events such as protection trips and sampled values with comparatively slower telemetry streams like AMI and SCADA signals, as well as broader market data. The architecture must therefore support mixed criticality and multirate time bases, ensuring seamless responsiveness without compromising stability. The second design choice involves semantic interoperability, a prerequisite for avoiding brittle point-to-point integrations that hinder scalability and adaptability. Here, data modeling frameworks ranging from AutomationML, which codifies asset structures, to companion specifications and CIM-style semantics

play a pivotal role in encoding topology, ratings, and operational state. These standards allow diverse simulators, optimization engines, and analytical tools to exchange context-rich messages without the need for bespoke adapters, thus fostering a robust and modular ecosystem (Kamrul & Omar, 2022; Schroeder et al., 2016). The third design choice reflects the inherently cyber-physical nature of grids, which span multiple disciplines and domains. To capture this complexity, most DT stacks rely on co-simulation frameworks capable of orchestrating interactions among electromagnetic transients, phasor-domain dynamics, communication latencies, and even adjacent domains such as building energy models or electric vehicle charging systems (Palensky et al., 2017). These architectural decisions are never merely technical; rather, they set the boundaries of what an operator can reliably ask of a twin, whether that entails evaluating Volt/VAR control setpoints, forecasting congestion, accelerating restoration timelines, or conducting sensitivity tests for DER dispatch, all while ensuring traceability of inputs, model versions, and decision pathways across organizational boundaries.

Figure 3: Digital Twin Concepts and Reference Architectures for Smart Grids



Several domain-specific architectural blueprints have emerged that crystallize these guiding principles and adapt them effectively for electric networks, offering structured pathways toward practical deployment. One of the most widely cited manufacturing DT reference models emphasizes a clear separation of concerns between data, model, and service layers while formalizing practices of model management, including versioning, calibration, and uncertainty propagation, all of which translate seamlessly into grid contexts where telemetered states must be reconciled with solved power flows and data-driven surrogates (Lu et al., 2020; Kamrul & Tarek, 2022). From a systems-engineering perspective, scholars further advocate for the explicit integration of a “digital thread,” which operates as a connective tissue binding requirements, asset hierarchies, and operational decisions together; within grid twins this thread becomes especially vital, as it allows topology changes, parameter updates, and operator interventions to be systematically traced back to validated sources, ensuring both consistency and accountability in dynamic environments (Madni et al., 2019; Mubashir & Abdul, 2022). Complementing these theoretical constructs, proof-of-concept demonstrations have underscored the importance of viewing DTs as part of an active ecosystem rather than as static replicas. Such studies illustrate operational loops in which twins continuously ingest live data streams, perform rapid what-if analyses, and disseminate tailored recommendations to human-in-the-loop decision-making platforms, thereby bridging automated intelligence with operator oversight (Haag & Anderl, 2018). Moreover, smart-grid-specific models like the OKDD (ontology-body, knowledge-body, data-body, digital-portal) framework elevate knowledge and ontology layers to first-class elements, ensuring that equipment semantics, procedural rules, and health standards are encoded alongside raw data schemas. This layered ontological integration facilitates reusability and interoperability across multiple scales, from unit-level twins to complex

system-of-systems applications encompassing substations and distribution networks (Jiang et al., 2022). Collectively, these architectural paradigms form the structural and procedural scaffolding necessary for AI-enabled DTs to deliver robust energy-efficiency services such as loss minimization, Volt/VAR optimization, and predictive maintenance with verifiable alignment to measurements, validated models, and governance mechanisms (Negri et al., 2017; Muhammad & Kamrul, 2022).

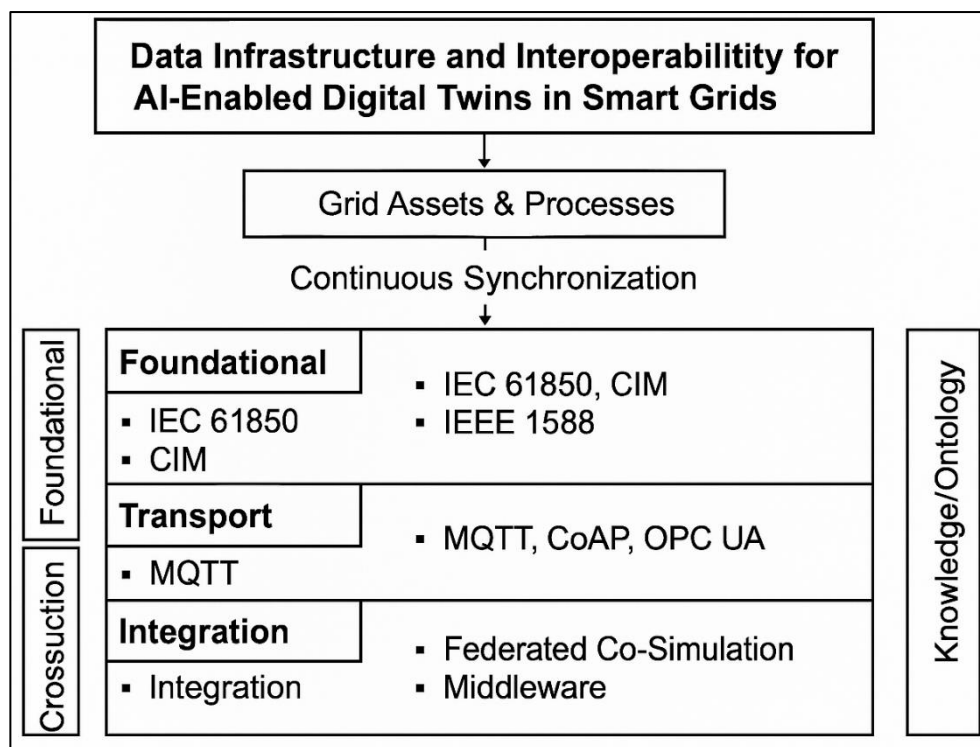
AI-Enabled Digital Twins in Smart Grids

A robust data infrastructure forms the essential connective tissue that allows AI-enabled digital twins (DTs) to ingest, harmonize, and reason over heterogeneous operational technology (OT) and information technology (IT) data streams, ensuring that the twin remains an accurate and actionable representation of the grid. Foundational to this infrastructure are domain-specific standards that encode semantics and enforce structured exchange rules, which enable interoperability across diverse devices and systems. Within substations and field equipment, IEC 61850 defines both the data models for assets and the engineering and operational workflows, establishing decades of deployment evidence that underscore its role as a linchpin for device-level interoperability that DTs can reliably exploit (Ayello & Lopes, 2023; Reduanul & MohShoeb, 2022). Yet semantic standardization alone is insufficient to achieve seamless integration; converged messaging stacks are required to bridge substation, control center, and enterprise analytics layers. Studies demonstrate that integrating IEC 61850's Substation Configuration Language (SCL) with OPC UA's service-oriented information model substantially reduces integration friction by unifying modeling and transport abstractions, a design pattern directly applicable to DT ingestion layers (Cavaliere & Regalbuto, 2016). Time-critical DT functions, such as protection-aware state estimation or event-triggered model retraining, depend on sub-microsecond alignment across publishers, brokers, and consumers, and empirical evaluations of IEEE-1588/PTP in IEC 61850 environments provide guidance on the accuracy and limitations of commercially available clocks and switches, informing how tightly DT pipelines can be synchronized without dedicated timing cabling (Han et al., 2019; Sabuj Kumar & Zobayer, 2022). Finally, at the model and integration tier, software engineering practices emphasize explicit requirements for standard-conformant model exchange and co-simulation, for example using FMI-based interchange and profile-driven validation, enabling DTs to couple physics-based simulators, telemetry adapters, and optimization services without reliance on brittle, bespoke adapters (Gómez et al., 2020). This combination of standards, messaging, timing, and model governance forms the backbone of reliable, high-fidelity DT operations in complex energy systems. Beyond substation boundaries, DT data fabrics increasingly span edge–cloud hierarchies. Reviews of “smart grid meets edge computing” show how edge orchestration reduces backhaul load, supports localized analytics, and enables hierarchical control capabilities DTs exploit for scalable feature extraction, online learning, and closed-loop actuation (Chen et al., 2021). In parallel, “edge intelligence” surveys detail patterns for partitioning model training/inference between constrained gateways and cloud, codifying design choices for DT components such as anomaly detectors, forecasting models, and reinforcement learning controllers deployed near assets (Onumanyi et al., 2022; Sadia & Shaiful, 2022). Since DTs depend on high-fidelity, reliable streams, application-layer protocol choices matter: comparative experiments show tradeoffs among MQTT (brokered, TCP), CoAP (RESTful, UDP), and OPC UA (rich semantics, session-oriented), with latency, jitter stability, and packet loss sensitivities varying by topology and traffic intensity evidence that DT data buses should be protocol-polyglot and topology-aware rather than protocol-monolithic (Seoane et al., 2021). Security overlays (e.g., DTLS/OSCORE) further alter performance; measurements of secure CoAP versus secure MQTT quantify overheads and can guide where to terminate crypto (edge vs. broker) so DTs preserve both confidentiality and real-time guarantees (Laaroussi & Novo, 2021; Sazzad & Islam, 2022). Practically, this means DT pipelines adopt tiered brokers, event-stream processors, and time-series stores close to sources, apply schema governance and versioned contracts at the edge, and use service meshes in the cloud while continuously validating end-to-end service-level objectives against empirical protocol behavior reported in the literature.

Interoperability at the system-of-systems level is a critical enabler for AI-driven digital twins, allowing them to integrate multiple domains such as power, communications, and markets into coherent, actionable operational views. Achieving this level of integration depends on co-simulation frameworks and federated middleware that preserve synchronization and causality across heterogeneous components while blending real-time measurements with simulated counterfactuals. Federated frameworks, such as the Federation of Networked Control Systems

(FNCS), have demonstrated repeatable and robust coupling between power-system simulators and communication models, providing a template for scalable, domain-spanning DT architectures (Huang et al., 2017; Noor & Momena, 2022). Recent research in high-DER environments extends these principles, illustrating that distributed co-simulation can capture multi-rate and multi-fidelity dynamics that monolithic solvers often overlook. This capability is particularly valuable for DT validation and test harnesses, enabling the verification of AI model behavior under contingencies, rare events, or extreme conditions prior to deployment in live networks (Chagas & Tomim, 2022; Akter & Razzak, 2022). These findings collectively underscore the importance of an interoperability stack composed of layered capabilities. At the foundational layer, standards-based semantic models such as IEC 61850 and CIM, combined with precise synchronization using IEEE-1588, ensure consistent and meaningful ingestion of asset states. At the transport layer, protocol-aware, edge-anchored data movement using MQTT, CoAP, or OPC UA ensures reliable communication with measured security and quality-of-service guarantees. At the integration layer, federated co-simulation and middleware maintain cross-domain consistency, facilitating holistic operational insights. When such a stack is implemented alongside rigorous requirements management and conformance testing, it produces a DT data backbone that is portable across vendors, adaptable to evolving grid infrastructures, and sufficiently robust to host safety-critical AI-driven services. This layered approach ensures that digital twins can act as trustworthy, extensible platforms for monitoring, optimization, and control across complex, modern energy networks.

Figure 4: Data Infrastructure and Interoperability Framework for AI-Enabled Digital Twins



AI for Forecasting to Drive Efficiency

Forecasting functions as the computational “look-ahead” that renders smart-grid digital twins (DTs) actionable, transforming raw telemetry into anticipatory control for energy efficiency. Within a DT, the analytics stack integrates high-resolution state estimation with predictive models that project near-term load profiles, renewable generation outputs, and market prices, enabling the twin to schedule assets and adjust control setpoints to minimize losses, curtailment, and operational waste. Artificial intelligence (AI) techniques, particularly deep learning architectures, are central to this capability, as they can discern complex nonlinear patterns from streams of smart-meter data, weather forecasts, and contextual operational signals. Probabilistic load forecasting extends the value of point predictions by generating distributions, allowing DT optimization layers to hedge

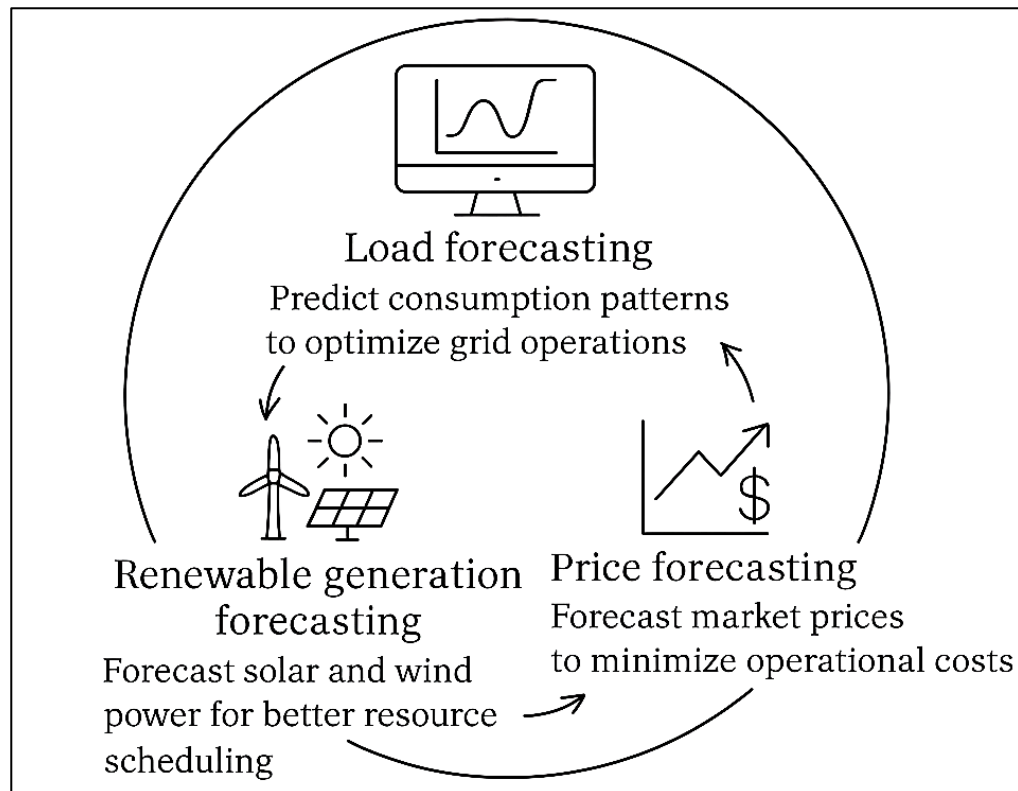
against uncertainty, reduce reserve margins, and maintain reliability without overprovisioning (Hong & Fan, 2016). At both feeder and household scales, recurrent neural networks, including long short-term memory (LSTM) models and pooling deep RNNs, consistently outperform traditional regression approaches for short-term load prediction, effectively capturing volatile end-use behaviors that influence peak demand, transformer stress, and localized losses (Kong et al., 2017; Shi et al., 2018). Embedding these forecasts within supervisory control routines such as model predictive controllers enables tangible operational benefits, including flattened demand profiles, decreased distribution losses, and lower demand charges through pre-cooling, storage pre-charging, or demand-shifting strategies. Smart-meter analytics further enhance these outcomes by providing fine-grained, user-specific baselines and anomaly detection, which refine forecast accuracy and improve the DT's decision-making fidelity (Adar & Md, 2023; Yildiz et al., 2017). Collectively, the coupling of probabilistic and deep learning forecasts equips DTs with foresight, allowing them to strategically trade minor forecast uncertainty against measurable efficiency gains, including reductions in curtailment, cycling losses, and idle generation, thereby converting streaming data into predictive, actionable, and economically valuable operational intelligence.

A second critical lever for energy efficiency in smart-grid digital twins (DTs) arises from the accurate forecasting of variable renewable generation, particularly photovoltaic (PV) and wind power, whose outputs are inherently weather-dependent and uncertain. Systematic reviews indicate that hybrid forecasting pipelines, which integrate satellite imagery, sky-imager features, numerical weather predictions, and machine-learning post-processing, significantly enhance the accuracy of PV power predictions across intraday horizons, enabling more reliable scheduling and dispatch (Antonanzas et al., 2016; Golam Qibria & Hossen, 2023). Similarly, wind-power forecasting has benefited from ensemble approaches and machine-learning corrections to traditional numerical weather models, producing measurable operational gains in unit commitment, reserve allocation, and curtailment reduction (Foley et al., 2012). Advancements in transformer-based neural network architectures tailored for energy data have further improved short-term forecasting accuracy and computational efficiency, supporting real-time control loops within DTs and reducing latency in decision-making (Capretz et al., 2022). Beyond purely data-driven approaches, the integration of physics-informed learning, such as physics-informed neural networks (PINNs), allows DTs to embed fundamental grid constraints, conservation laws, and device equations directly into the training objectives of predictive models. This approach ensures that forecasted renewable outputs and simulated network states remain consistent with the physical realities of the grid, enhancing robustness under sparse, noisy, or outlier-laden observations (Istiaque et al., 2023; Raissi et al., 2019). By combining weather-aware PV and wind forecasts with physics-consistent load evolution, a DT can strategically schedule energy storage, flexible loads, and distributed resources with reduced safety margins. This precision translates into concrete efficiency outcomes, including lower spinning reserves, decreased ramping requirements, and minimized clipping of renewable generation, all of which are observable in the twin's power-flow simulations and thermal analyses. Collectively, these forecasting advancements allow DTs to convert uncertain renewable generation into actionable, optimized operational strategies that enhance reliability and economic performance while reducing losses across the distribution network.

Furthermore, forecasting price signals and market conditions provides a crucial bridge between technical efficiency and economic optimization in smart-grid digital twins (DTs). When a DT co-optimizes energy use, emissions, and operational costs, accurate day-ahead and intraday price predictions enable strategic adjustments of flexible loads, HVAC setpoints, and energy storage dispatch, thereby maximizing value per kilowatt-hour while maintaining system reliability. Empirical benchmarks evaluating twenty-seven algorithms across multiple electricity markets demonstrate that deep learning approaches particularly recurrent and convolutional architectures perform competitively in day-ahead price forecasting, supplying actionable signals that DT controllers can incorporate into automated decision-making loops (Lago et al., 2018; Akter, 2023). Foundational reviews in electricity price forecasting further emphasize the importance of rigorous feature engineering, including calendar, weather, and demand effects, as well as robust evaluation and ensemble modeling techniques, all of which remain essential when price forecasts inform real-time operational actions within a DT framework (Hasan et al., 2023; Weron, 2014). Methodologically, these forecasts complement probabilistic load and renewable generation predictions: probabilistic load forecasts safeguard operational reliability, renewable forecasts enable flexibility and anticipatory

scheduling, and price forecasts monetize optimal timing, creating a holistic foundation for both technical and economic efficiency. By embedding these predictive layers into a continuously updated DT, operators gain a forward-looking perspective that integrates expected demand, renewable availability, and market signals, allowing assets to be scheduled proactively rather than reactively. The result is a system capable of achieving measurable reductions in energy losses, optimized dispatch of distributed energy resources, improved utilization of storage, and cost savings through arbitrage opportunities. In this way, AI-driven price forecasting, when fused with load and renewable predictions within a unified digital twin, transforms operational decision-making, providing a structured, data-rich pathway to maximize energy efficiency, reliability, and financial performance across modern power distribution networks.

Figure 5: AI-Enabled Forecasting Functions in Digital Twins for Energy Efficiency



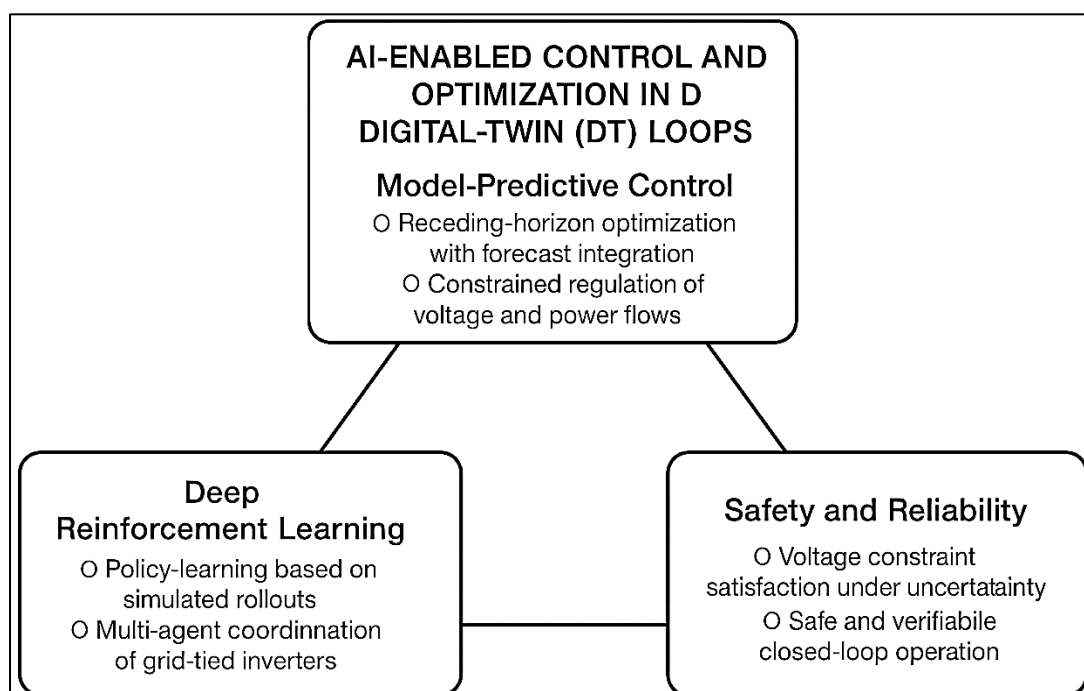
AI-Enabled Control and Optimization in Digital-Twin (DT) Loops

Within digital-twin (DT) closed-loop frameworks for smart grids, artificial intelligence (AI) plays a transformative role in coordinating both slow, discrete actuators such as on-load tap changers (OLTCs) and capacitor banks, alongside fast, continuous resources including inverter VARs and battery storage systems. Central to this coordination is model-predictive control (MPC), a long-established method that integrates forecasts into constrained, receding-horizon optimization, allowing the DT's state estimator and simulator to propose feasible control trajectories that maintain voltage within prescribed limits, respect equipment cycling constraints, and minimize network losses. Distribution-level voltage regulation studies exemplify MPC's capacity to enable a DT to "preview" future operating states and iteratively correct setpoints as real-time measurements arrive, thereby enhancing both stability and responsiveness (Masud et al., 2023; Valverde & Cutsem, 2013). Enhancements that embed actuator physics explicitly accounting for dead-bands, discrete step changes, and inherent delays in OLTC operations reduce discrepancies between simulated and field behavior, mitigating the risk of control chatter or infeasible tap sequences (Colas et al., 2017). Event-triggered predictive Volt/VAR control further refines this approach by initiating optimization only when the DT identifies operating regimes that warrant intervention, such as rapid photovoltaic ramps, thereby reducing computational overhead while ensuring adherence to operational constraints (Singh, 2021). These MPC-based strategies constitute the interpretable, verifiable, and constraint-

aware “glass box” core of many DT loops, offering a structured foundation for decision-making. Increasingly, this core is augmented by AI learning agents capable of handling uncertainty, high-dimensional interactions, and nonlinearity in the system that traditional models struggle to capture. By integrating MPC with data-driven learning, DTs not only retain predictability and compliance with operational limits but also expand their capability to respond adaptively to evolving conditions, ultimately improving feeder-level energy efficiency, asset utilization, and overall system reliability within a rigorously validated, model-informed control environment.

Deep reinforcement learning (DRL) operates within the digital-twin (DT) framework as a policy-learning layer, enabling control strategies to be derived directly from simulated rollouts and streaming telemetry, thereby reducing reliance on hand-crafted rules. In both transmission and distribution networks, centrally trained yet decentrally executed multi-agent DRL architectures have demonstrated the ability to coordinate numerous inverter agents for autonomous voltage regulation. The DT serves as a safe, high-fidelity environment in which these agents can explore operating contingencies, test policies, and evaluate performance prior to field deployment, ensuring system reliability and compliance with operational constraints (Sultan et al., 2023; Wang et al., 2020). Multi-agent DRL formulations scale effectively to feeders with high photovoltaic penetration by leveraging spatial locality and sparse inter-agent communication, allowing each agent to act on local observations while sharing essential network information through the DT's virtual environment (Cao et al., 2021; Hossen et al., 2023). Beyond inverter-centric control, DRL has been extended to hybrid assets such as multi-terminal soft-open points (SOPs), which provide flexible routing of active and reactive power; here, learned policies can exploit network reconfigurability and uncertainty while the DT enforces constraints and verifies feasibility before any command reaches physical devices (Li, 2022). Two-timescale DRL implementations are particularly compatible with DTs: fast agents adjust inverter setpoints at sub-minute intervals using the DT's most recent state, whereas slower agents orchestrate discrete devices hourly, together achieving reductions in long-term voltage deviation, technical losses, and feeder stress, as validated in large-scale simulations (Tawfiqul, 2023; Yang et al., 2020). Embedding DRL within the DT loop permits rigorous stress-testing against stochastic weather and load scenarios, cyber-latency patterns, and topology changes, allowing operators to assess the robustness, safety, and efficiency of emergent policies before applying them in live grids. This integration of learning, simulation, and controlled experimentation ensures that DRL-driven DTs provide actionable, risk-mitigated, and adaptive energy management strategies across complex, high-renewable networks.

Figure 6: AI-Enabled Control and Optimization in Digital-Twin Loops for Smart Grids



Safety and reliability form the foundation of effective closed-loop control when integrating learning agents into digital-twin (DT) frameworks for smart grids. Topology-aware deep reinforcement learning (DRL) enhances policy performance by embedding graph structures that reflect electrical neighborhoods, enabling agents to exploit spatial correlations and local connectivity; in DT-in-the-loop evaluations, this approach improves voltage profiles and maintains robustness under partial observability, particularly during dynamic feeder reconfigurations (Xiang, 2023). Complementing this, safe DRL methods explicitly constrain closed-loop behavior through techniques such as Lyapunov-style conditions, stability-certified critics, or projection operators, ensuring that reactive-power interventions keep voltages within operational limits even in the presence of non-stationary photovoltaic generation and load fluctuations. The DT serves a critical role in these setups by generating counterfactual rollouts that certify policy safety and feasibility before any command is issued to the physical network (Cui et al., 2022; Shamima et al., 2023). Moreover, co-optimization strategies extend beyond voltage regulation to incorporate ancillary objectives such as loss minimization, device wear mitigation, and operational cost considerations; multi-agent safe DRL architectures, supported by high-fidelity DT states, allow network operators to simulate disturbances, score policies in “shadow mode,” and validate performance improvements prior to live deployment (Hossain et al., 2023; Ashraf & Ara, 2023). By combining model-predictive control’s constraint-handling rigor with DRL’s pattern-seeking adaptability, DT-enabled AI loops offer a balance of flexibility, safety, and verifiability. Every control action can be tested, monitored, and gated within the DT, creating a continuous feedback layer that ensures operational compliance and reduces the risk of inadvertent violations. In essence, this integrated framework transforms conventional voltage and reactive-power management into a proactive, intelligence-driven system, where learning, simulation, and validation coalesce to deliver resilient, verifiable, and efficient grid operations across complex, high-penetration renewable networks.

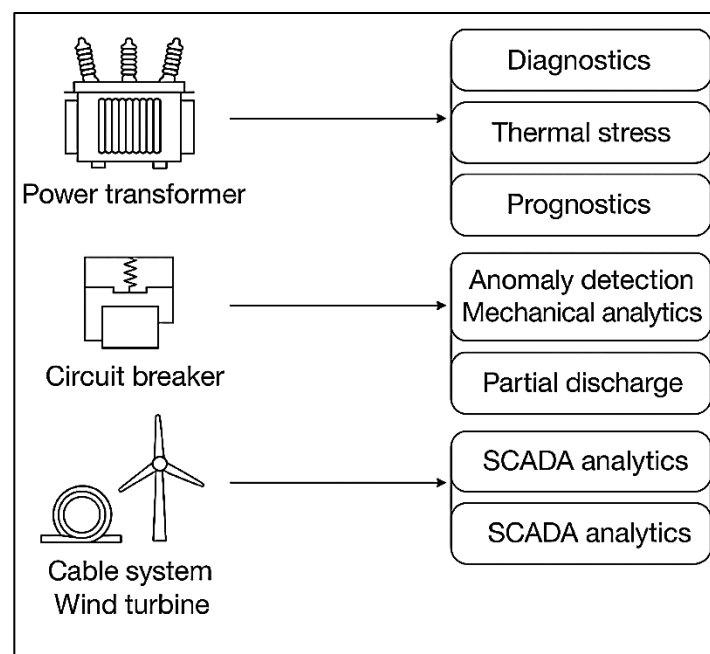
Predictive Maintenance and Asset-Health Twin Modeling

Digital twins (DTs) become particularly powerful when they are “asset-centric,” i.e., when the twin’s purpose is to continuously assess the health of a specific class of equipment and to anticipate failures with enough lead time for low-cost interventions. In power transformers, for example, decades of practice around dissolved gas analysis (DGA) provide a rich foundation of features that DTs can ingest alongside loading, ambient, and topology context to infer latent fault modes. Early work showed that statistical and machine-learning treatments of DGA can outperform rigid rule-bases by exploiting multivariate structure and nonlinearity, opening the door to twin-driven diagnostics that update as new samples arrive (Mirowski & LeCun, 2012). In parallel, thermal stress remains a primary ageing driver; here, asset-health twins commonly couple compact thermo-electrical surrogates with data-driven learners to estimate winding hot-spot temperatures under varying load and cooling states and to convert those estimates into loss-of-life metrics. Particle-filter-optimized support vector regression for dry-type units is a representative approach, achieving accurate hot-spot forecasts with relatively few parameters useful when detailed geometry is unavailable to the operator (Sun et al., 2021). Beyond single-task predictors, modern prognostics embed sequence models to forecast the remaining useful life (RUL) of assets from multichannel telemetry streams. Transformer-encoder architectures, for instance, learn long- and short-range temporal dependencies directly from condition-monitoring data and thus fit naturally into DTs that stream features from SCADA/PMU historians and on-board sensors, enabling the twin to present a probabilistic RUL with quantified uncertainty (Michau et al., 2021; Sanjai et al., 2023). Together these strands DGA-based diagnostics, thermal stress estimation, and sequence-learning prognostics form the backbone of transformer-centric health twins that translate continuous sensing into actionable maintenance schedules.

Circuit breakers (CBs) exemplify a complementary asset-centric paradigm in which digital twins (DTs) focus on electromechanical subsystems and the management of fast transients, illustrating how predictive maintenance can extend beyond conventional, calendar-based schedules. Many CB degradation mechanisms reveal themselves as subtle variations in auxiliary coil currents, contact travel profiles, or mechanism vibration signatures, making data-driven condition assessment particularly effective when combined with feature learning and multi-signal fusion techniques. For instance, back-propagation neural networks applied to heterogeneous CB records demonstrate how a DT can integrate asynchronous measurements to generate an evolving equipment-state score, thereby supporting risk-based maintenance decisions rather than relying solely on

predetermined inspection intervals (Geng & Wang, 2020). Recent advances further extend this paradigm from static condition evaluation to short-horizon prediction. Long short-term memory (LSTM) networks are employed to forecast coil-current trajectories and mechanical movements one step ahead, and these forecasts are subsequently converted into derived mechanical indicators such as speed and bounce before classification with support-vector machines. This methodology fits naturally into a DT loop, where predicted anomalies can trigger additional sensing actions or controlled test shots, enabling preemptive intervention without compromising operational continuity (Akter et al., 2023; Zheng et al., 2023). When high-frequency vibration data are available, deep-learning classifiers can achieve over 95 percent accuracy in identifying fault types for specific breaker families, providing a robust front-end for a DT that executes shadow diagnostics after each operation and updates the residual-life model of the asset (Chen et al., 2023). In practice, a CB-health twin synthesizes coil-current analytics, vibration-based inference, and mechanistic constraints, delivering interpretable, actionable alerts to maintenance personnel. These alerts highlight probable root causes and prioritize attention on the most degraded subassemblies, allowing maintenance teams to make informed decisions that maximize operational reliability, minimize downtime, and preserve system safety.

Figure 7: Predictive Maintenance and Asset-Health Digital Twin Modeling in Smart Grids



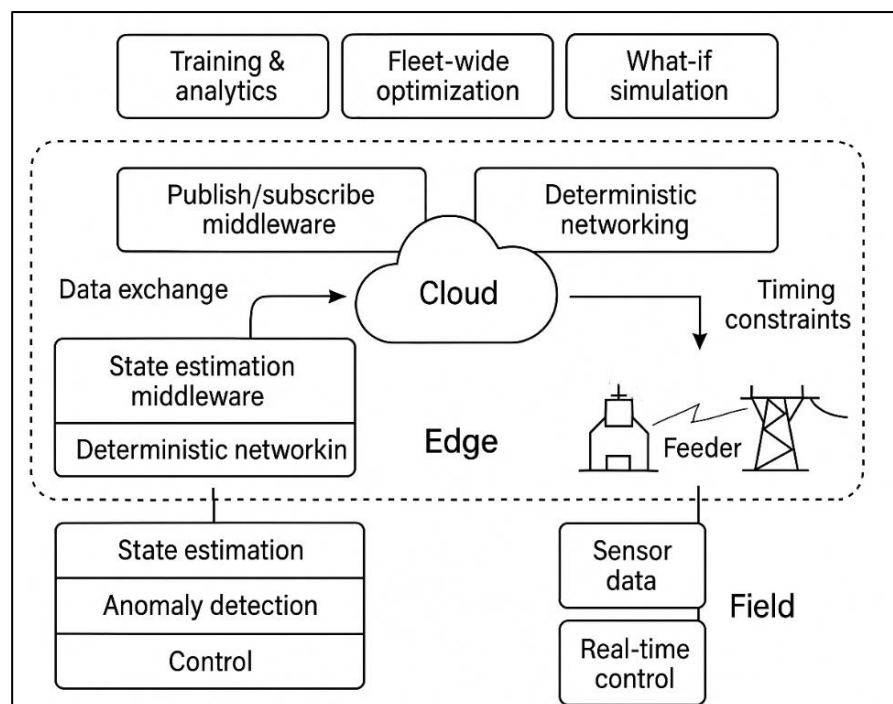
Cable systems and rotating renewable assets extend the asset-centric scope of digital twins (DTs) by illustrating how twin frameworks generalize across distinct physical domains while maintaining predictive and prescriptive capabilities. In high-voltage cable networks and connectors, partial discharge (PD) is a well-established precursor to insulation degradation and eventual failure, and deep convolutional neural networks trained on raw or minimally processed waveform data have demonstrated robust PD versus no-PD discrimination. These models retain interpretability through saliency-style pulse activation maps, enabling the DT not only to detect anomalies but also to explain the basis for decisions, which is essential for regulated environments and operational trust (Michau et al., 2021). Modern field deployments increasingly employ multimodal sensing, such as high-frequency current transformers, ultrasonic probes, and ultra-high-frequency detectors, and multi-sensor fusion models within the DT enhance robustness to noise, nonstationary signals, and complex termination geometries typical in traction or substation cabling (Razzak et al., 2024; Li et al., 2024). For wind turbines, asset-health twins typically leverage SCADA data streams, and systematic reviews demonstrate that effective signal preprocessing, feature engineering, and learning architectures converge on reliable indicators for gearbox, generator, and blade faults (Tautz-Weinert & Watson, 2017). Hybrid models that integrate SCADA telemetry with vibration signatures further improve

sensitivity to incipient failures, enabling predictive scheduling of inspections, maintenance interventions, and spare-part logistics before performance degradation or catastrophic events occur (Turnbull et al., 2019). Across these heterogeneous assets, a consistent operational paradigm emerges: the DT ingests diverse, asset-specific data, applies trained predictive models to estimate condition, anomaly likelihood, and remaining useful life, and outputs interpretable and verifiable recommendations that directly inform maintenance actions (Mo et al., 2021). Crucially, these predictions are anchored in the physics and operational constraints of the underlying equipment, ensuring that the DT preserves fidelity to real-world dynamics while delivering actionable intelligence for condition-based and predictive maintenance strategies.

Edge–Cloud Deployment and Real-Time Constraints

Designing AI-enabled digital twins that actually meet the hard timing budgets of smart-grid operations starts with where computation lives along the cloud-to-edge continuum. The edge model pushing analytics and control closer to data sources reduces transport delay, alleviates backbone congestion, and supports context-aware decision making under volatile network conditions (Istiaque et al., 2024; Satyanarayanan, 2017; Shi et al., 2016). In parallel, fog/edge layers interposed between field devices and centralized clouds provide intermediate compute, storage, and networking primitives that can be orchestrated to satisfy latency and reliability targets without overprovisioning the core (Bonomi et al., 2012). Within this layered topology, latency is not a single number but the sum of sensor sampling, serialization, queueing, transport, inference, and actuation delays each influenced by placement, contention, and scheduling. Latency-aware application management strategies therefore decompose twin workloads (state estimation, anomaly detection, control policy evaluation) into modules whose placement is co-optimized for proximity, compute capacity, and update frequency to keep loop-closure within protection and control deadlines (Mahmud et al., 2018; Akter & Shaiful, 2024). In field deployments, container-orchestrated microservices are attractive for modularity and portability, but the marginal overheads they introduce especially for I/O-heavy paths must be budgeted explicitly when closing real-time loops (Santos et al., 2018). The architectural implication is that digital-twin pipelines should be partitioned so that time-critical inference and control run on substations or feeders, while history-rich training, fleet-wide optimization, and what-if simulation remain in the cloud, with well-defined service-level objectives (SLOs) that reflect end-to-end timing constraints (Khattach et al., 2025).

Figure 8: Edge–Cloud Deployment Strategies and Real-Time Constraints in AI-Enabled Digital Twins



Meeting service-level objectives (SLOs) at scale in digital twin (DT) deployments relies heavily on deterministic networking to ensure that critical data flows remain timely, reliable, and predictable across distributed twin components. Time-Sensitive Networking (TSN) enhances standard Ethernet with traffic shaping, precise time synchronization, and bounded-latency scheduling, allowing high-priority flows such as phasor updates, topology changes, and control set-points to traverse shared networks with minimal jitter and guaranteed delay bounds (Seol et al., 2021). Building atop this deterministic transport layer, publish/subscribe middleware frameworks, notably OPC UA integrated with TSN, provide vendor-neutral data modeling and deterministic message delivery, enabling DT microservices to selectively subscribe to latency-sensitive streams while relegating bulk telemetry and lower-priority updates to best-effort classes (Hasan et al., 2024; Trifonov & Heffernan, 2023). Even with these foundational mechanisms, practical edge–cloud coordination must accommodate intermittent connectivity, asymmetric resources, and the operational realities of distributed grid environments. Platforms such as KubeEdge extend Kubernetes semantics to edge nodes, supporting offline-tolerant metadata synchronization, workload orchestration, and seamless state reconciliation upon reconnection, ensuring that critical control processes continue functioning safely during cloud partitions or network backhaul degradation (Xiong et al., 2018). From an operational planning perspective, deployment strategies must carefully align control-theory deadlines, including tens-of-milliseconds requirements for feeder protection and sub-second windows for remedial action schemes, with the guarantees provided by TSN traffic classes, local scheduling policies, and admission-control mechanisms. This alignment prevents high-criticality components from being starved by background analytics, batch telemetry transfers, or lower-priority services, preserving both the timeliness and determinism required for safe, reliable, and efficient grid operation. By combining deterministic networking, middleware abstraction, and edge-aware orchestration, DTs can maintain synchronized, high-fidelity operational models that support real-time decision-making, fault response, and closed-loop control across complex, multi-layered smart-grid infrastructures.

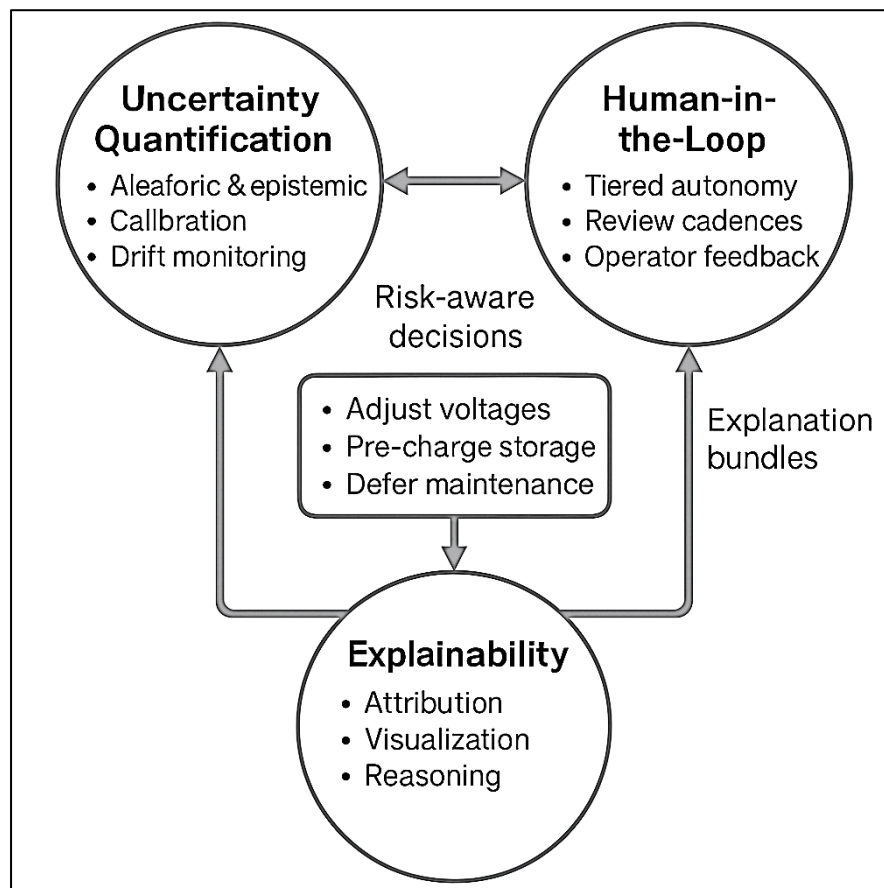
Achieving sustained real-time performance in digital twin (DT) deployments demands meticulous attention to system-level trade-offs highlighted by module placement and orchestration. Latency-aware management demonstrates that relocating a subset of operators such as feature extraction, threshold evaluation, or local policy inference to fog or edge nodes can substantially reduce end-to-end delays below protective thresholds while also alleviating backhaul congestion, provided that placement strategies account for input data rates, directed acyclic graph (DAG) dependencies, and burstiness patterns in telemetry streams (Mahmud et al., 2018). Empirical studies further reveal that naively containerizing every processing stage may increase execution time and energy consumption, particularly for I/O-bound workloads, which can compress timing margins or violate latency constraints; performance profiling therefore guides a hybrid approach in which bare-metal execution is reserved for the most latency-sensitive paths while containerized services handle ancillary or compute-flexible functions (Santos et al., 2018). At the pipeline level, stream-processing frameworks such as Kafka and Spark provide scalable ingestion, micro-batched processing, and near-real-time analytics, yet their checkpointing and buffering semantics must be carefully tuned to the twin's freshness requirements and grid control horizons, with stateful operators preferentially pinned to edge locations to minimize age-of-information (Aol) under transient conditions (Khattach et al., 2025). Architecturally, these considerations converge into a coordinated edge–cloud fabric: deterministic network links and priority-aware queues transport time-critical topics to latency-hardened services at the edge, resilient orchestration manages drift, faults, and reconvergence, and the cloud accommodates compute-intensive learning, historical analytics, and fleet-level orchestration. The entire deployment is governed by explicit policies that encode timing constraints, reliability targets, and safety margins, ensuring that the DT can operate in real time across a geographically distributed, heterogeneous infrastructure while preserving responsiveness, fidelity, and operational security.

Uncertainty, Explainability, and Human-in-the-Loop

Uncertainty quantification (UQ) is the backbone of trustworthy AI inside digital-twin loops because efficiency decisions whether to pre-charge storage, adjust Volt/VAR setpoints, or defer maintenance hinge on how confidently forecasts and state estimates are believed. A practical decomposition distinguishes aleatoric uncertainty (data noise intrinsic to demand, weather, or sensors) from epistemic uncertainty (model and data-coverage limits), each demanding different treatments for estimation and for downstream optimization. Modern reviews detail a toolkit that spans Bayesian

and non-Bayesian approaches, including ensemble methods, variational families, stochastic regularization, and calibration post-processing, and emphasize that credible intervals must be reported and stress-tested alongside point predictions if models are to be operated in closed loop (Abdar et al., 2021). In grid-efficiency contexts, credible uncertainty is not merely decorative: scheduling and control modules inside the twin apply risk-aware objectives that reward sharp yet reliable distributions rather than overconfident point forecasts. Proper scoring rules offer rigorous criteria for this trade-off; continuous ranked probability score (CRPS) and log score remain standard because they are strictly proper encouraging truthful distributions and penalizing miscalibration (Gneiting & Raftery, 2007; Tawfiqul et al., 2024). When data are nonstationary or heterogeneously sampled, distribution-free methods such as conformal prediction can still provide finite-sample coverage guarantees, enabling the twin to expose prediction intervals with user-chosen confidence levels without assuming parametric forms a practical advantage for operators who must set explicit safety margins (Angelopoulos & Bates, 2023; Rajesh et al., 2024). Finally, because smart-grid data are subject to seasonal shifts, topology changes, and evolving DER fleets, drift is the norm rather than the exception; a broad survey of concept-drift handling shows that effective monitoring combines windowed detectors, adaptive learners, and selective retraining so that uncertainty estimates remain meaningful as the data-generating process changes functionality that a digital twin can institutionalize as part of its data-quality and model-lifecycle governance (Gama et al., 2014).

Figure 9: Human-in-the-Loop Integration in AI-Enabled Digital Twins



Explainability methods complement UQ by translating complex model behavior into evidence that is intelligible to grid engineers and operators who must validate recommendations before acting. Comprehensive taxonomies separate global (model-level) from local (instance-level) explanations and catalog families such as surrogate models, attribution/importance scores, counterfactuals, and example-based reasoning, together with desiderata like faithfulness, stability, and comprehensibility (Guidotti et al., 2018). In practice, local explainers are frequently embedded at the twin's decision surface: LIME fits sparse, interpretable surrogates in the neighborhood of each query to expose which

features e.g., feeder loading percentiles, PV ramp indicators, or nodal voltages most influenced a recommended action (Ribeiro et al., 2016). SHAP provides an alternative grounded in cooperative-game theory; by attributing predictions to features via Shapley values and aggregating them to partial importance and interaction plots, operators gain consistent local and global narratives of why a controller favored one schedule over another (Lundberg et al., 2020). At the systems level, explainability is also a governance instrument: a widely cited review argues that responsible AI requires transparency about data provenance, model structure, and uncertainty, and that explanation artifacts should be auditable and aligned to human decision-rights rather than offered as post-hoc rationalizations (Arrieta et al., 2020; Subrato & Md, 2024). In a digital-twin setting, these tools can be operationalized as immutable “explanation bundles” attached to each optimization or control issuance, containing attribution vectors, counterfactual checks, and uncertainty summaries. Such bundles allow engineering review, ex-post measurement and verification, and incident analysis, and they provide training material for new operators closing the loop between model developers and field personnel while maintaining traceability across twin versions and data updates. Human-in-the-loop (HITL) design turns UQ and explainability into operational advantages by structuring how people and automation share tasks, attention, and authority during routine operations and disturbances. Foundational human-factors work models trust as a dynamic calibration problem: too little trust leads to wasted human labor and under-utilization of automation; too much trust induces automation bias and unsafe reliance, especially under distribution shifts that silently invalidate model assumptions (Lee & See, 2004; Ashiqur et al., 2025). Situation-awareness research likewise emphasizes that operators must perceive, comprehend, and project system state; digital-twin interfaces that externalize model beliefs, uncertainty ranges, and “what-if” outcomes directly support these levels and mitigate out-of-the-loop performance decrements (Endsley, 1995). Effective HITL patterns therefore codify *when* the twin acts autonomously, *when* it requests a human check, and *what evidence* must accompany that request. In practice, this maps to tiered autonomy with guardrails for example, the twin is permitted to dispatch reactive power within narrow bounds when uncertainty is low, but must seek human confirmation when forecast dispersion or drift indicators cross thresholds, or when recommended actions violate learned operator preferences. The workflow also defines review cadences for explanation bundles and drift dashboards, and prescribes how operator feedback is captured as labels, constraints, or counterexamples that refine future models. Together with disciplined uncertainty reporting (sharp but calibrated) and faithful explanations (locally accurate, globally coherent), a HITL twin promotes resilient, accountable efficiency improvements that can be audited, taught, and continuously improved without eroding operator authority or safety culture.

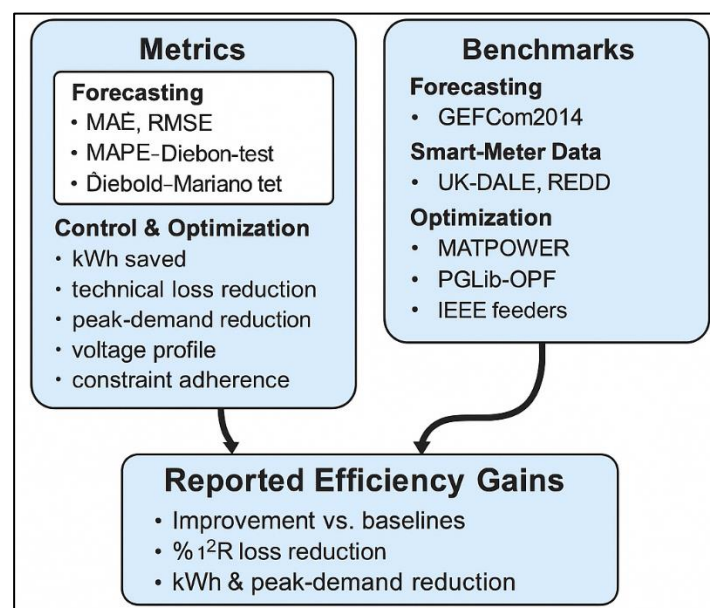
Metrics, Benchmarks, and Reported Efficiency Gains

A consistent evaluation grammar is essential to compare efficiency-oriented results across AI-enabled digital twins (DTs) in smart grids. At the forecasting layer that feeds most twin decisions, studies commonly report point-error metrics such as mean absolute error (MAE), root mean squared error (RMSE), and mean absolute percentage error (MAPE); while each emphasizes a different loss geometry, their joint reporting helps separate bias from variance and penalization of large errors from overall fit (Hyndman & Koehler, 2006). Because DTs often combine multiple forecasters or compare alternative pipelines, statistical comparators are equally important: the Diebold–Mariano test remains the workhorse for assessing whether two competing predictive schemes have significantly different expected losses under a chosen scoring rule, making it suitable for model selection in operational twins (Diebold & Mariano, 1995; Md Hasan, 2025). Lessons from large-scale forecasting benchmarks generalize well to grid contexts: the M4 competition showed how robust baselines, pooled information across horizons, and hybrid designs can outperform bespoke, highly tuned models in many settings an insight that argues for transparent and competitive baselining before deploying a forecaster inside a twin loop (Makridakis et al., 2018). On the control and optimization side, efficiency is typically quantified through kWh saved relative to a baseline (pre/post or counterfactual), percentage change in technical losses on feeders, peak-demand indicators (e.g., max/quantile reductions), voltage-profile quality (e.g., absolute voltage deviation and statutory compliance rates), and equipment-usage surrogates such as tap operations or inverter reactive-power duty. Evaluations further track constraint adherence (voltage and thermal limits satisfied), computational latency to verify real-time feasibility, and stability of savings under realistic perturbations. When reported together with clear baselines and uncertainty summaries, these

metrics allow reproducible claims about energy-efficiency improvements attributable to the DT and its embedded AI.

Benchmark datasets and test systems anchor comparability across studies by standardizing inputs, contexts, and scoring protocols. For forecasting and price-sensitive scheduling, the Global Energy Forecasting Competition 2014 (GEFCom2014) supplied open load, wind, solar, and price tracks with harmonized scoring, enabling repeatable assessment of probabilistic and point forecasts that many twins adopt wholesale or emulate in internal evaluations (Hong et al., 2016). For solar, comprehensive reviews confirm canonical feature sets and horizon-dependent evaluation practices sky-imager and satellite features, numerical weather prediction, and post-processing clarifying how accuracy should be judged at intra-hour to day-ahead horizons relevant to DT-informed scheduling (Inman et al., 2013). At the demand side, public smart-meter corpora support household and feeder-level analytics relevant to baselining and anomaly detection; the REDD dataset established a widely used appliance- and circuit-level benchmark for disaggregation and consumption modeling, while the UK-DALE corpus extended long-horizon, high-frequency monitoring across multiple homes both commonly reused to vet forecasting and detection modules before twin integration (Kelly & Knottenbelt, 2015; Kolter & Johnson, 2011; Sultan et al., 2025). Grid-side optimization and loss analysis typically rely on open test cases and solvers: MATPOWER popularized a transparent implementation of AC and DC power-flow/OPF problems with shareable cases, facilitating like-for-like comparisons of algorithms and permitting reproducible integration with DT simulators (Zimmerman et al., 2011). More recently, the PGLib-OPF collection curated feasible AC-OPF instances with documented difficulty and validation checks, offering a stronger basis for benchmarking optimization routines that a DT might invoke for Volt/VAR or reconfiguration studies (Coffrin et al., 2018; Sanjai et al., 2025). Foundational distribution-feeder cases, such as those introduced in classic reconfiguration work, remain the de facto stage for reporting loss reductions and voltage improvements under control or topology changes, providing a shared reference for DT-in-the-loop evaluations (Baran & Wu, 1989).

Figure 10: Metrics, Benchmarks, and Reported Efficiency Gains



When metrics and benchmarks are applied collectively, the literature generally reports energy-efficiency gains as measurable improvements relative to transparent baselines under clearly declared operating assumptions. Digital twins (DTs) orchestrating feeder-level interventions such as Volt/VAR optimization, topology reconfiguration, and storage dispatch often quantify outcomes in terms of percentage reductions in I^2R losses across feeders, accompanied by enhanced voltage-compliance statistics observed on standard test feeders, with experiments repeated across diverse loading conditions to demonstrate robustness and reliability (Baran & Wu, 1989; Coffrin et al., 2018; Hyndman & Koehler, 2006). DTs that embed forecasting within operational loops report kilowatt-hour

savings and peak-demand reductions, derived from counterfactual simulations or pre- and post-intervention analyses, with the credibility of these results reinforced by rigorous error metrics and significance testing aligned with competition-grade evaluation protocols, ensuring that efficiency gains reflect the twin's predictive capabilities rather than arbitrary assumptions (Hyndman & Koehler, 2006; Makridakis et al., 2018; Zimmerman et al., 2011). On the demand side, household- and feeder-level modules are frequently validated against public datasets, demonstrating that techniques such as load disaggregation, anomaly filtering, or augmented baselines improve downstream scheduling decisions, thereby linking open-benchmark performance directly to field-level key performance indicators (Inman et al., 2013; Kolter & Johnson, 2011). Across these multiple strands, best practices emphasize pairing quantitative efficiency outcomes with thorough disclosure of evaluation design, including baseline definitions, horizon and granularity of scoring, penalties for constraint violations, and timing budgets. This disciplined approach ensures that results are portable across utilities and vendors, enabling peer reviewers and practitioners to distinguish genuine energy savings produced by AI-enabled DT interventions from those arising merely from favorable data quality, network conditions, or controllability assumptions, thereby advancing both reproducibility and operational confidence in smart-grid twin deployments.

METHOD

This study followed the Preferred Reporting Items for Systematic Reviews and Meta-Analyses (PRISMA 2020) framework to ensure a systematic, transparent, and reproducible process, culminating in a final evidence base of 103 peer-reviewed articles. A prospective protocol was drafted and time-stamped prior to any searches, defining objectives, research questions, eligibility criteria, screening flow, data items, risk-of-bias domains, and synthesis plans; any subsequent deviations were recorded with justification. Studies were eligible if they were English-language, DOI-indexed journal or full conference papers explicitly addressing artificial-intelligence-enabled digital twins in smart-grid contexts and reporting energy-efficiency-relevant outcomes with sufficient methodological detail; non-archival items and papers without verifiable outcomes were excluded. Comprehensive queries were executed across major scholarly databases, complemented by forward-backward citation chasing; database-specific Boolean strings combined controlled vocabulary and keywords for "digital twin," "smart grid/microgrid/distribution network," "energy efficiency/loss reduction/Volt-VAR/demand response/optimal power flow," and "machine learning/deep learning/reinforcement learning/graph learning/physics-informed." Records were de-duplicated and independently screened in two stages (title-abstract, then full text) by paired reviewers; disagreements were resolved by consensus, and full-text exclusions were documented to populate the PRISMA flow. Using a piloted codebook, paired reviewers independently extracted bibliographic details, grid context and twin role, AI methods, datasets and interoperability standards, calibration/validation procedures, efficiency metrics and baselines, quantitative results, compute/latency constraints, and deployment maturity; discrepancies were reconciled and inter-rater agreement monitored. Each included study underwent structured critical appraisal spanning internal validity, external validity, measurement validity, and reproducibility, with study-level risk-of-bias judgments recorded for sensitivity analyses. Given design and metric heterogeneity, evidence was integrated through narrative synthesis with structured tabulation; where feasible, outcomes were harmonized to comparable indicators (for example, percentage loss or peak reduction) and summarized with appropriate dispersion while respecting baseline definitions and horizons; the review adheres to PRISMA 2020 reporting and presents an auditable flow from identification to inclusion for the 103 studies.

Screening and Eligibility Assessment

Screening and eligibility assessment proceeded in two sequential stages designed to operationalize the predefined inclusion criteria and to minimize selection bias, ultimately yielding 103 studies for synthesis. After exporting all records from the target databases into a reference manager, exact and fuzzy duplicate detection was performed and verified manually to ensure that variant metadata (for example, differing conference and journal versions) were consolidated under a single canonical entry before screening. Title-abstract screening was then conducted independently by two reviewers against a pilot-tested form derived from the protocol, which required explicit linkage to artificial-intelligence-enabled digital twins in smart-grid contexts and at least one energy-efficiency-relevant outcome (such as technical loss reduction, peak shaving, Volt/VAR performance, curtailment minimization, or maintenance/asset-health savings), together with peer-reviewed status,

a registered DOI, and English language. Records that were clearly out of scope general IoT frameworks without twin coupling, purely theoretical AI work without power-system application, non-archival items, or studies lacking quantifiable outcomes were excluded at this stage with reasons captured in a screening log. Ambiguous abstracts were advanced to full-text review to avoid erroneous exclusions. Full texts of the remaining records were retrieved through library subscriptions or publisher portals; where access barriers existed, alternative legitimate sources (author-hosted versions, institutional repositories) were used to ensure comprehensive assessment. During full-text eligibility appraisal, the same two reviewers independently verified the presence of a digital-twin construct with bidirectional data coupling or continuous synchronization, an AI component used for forecasting, diagnosis, optimization, or control within the twin loop, and efficiency-oriented metrics reported against an explicit baseline under clearly described conditions. Studies were excluded if the twin was only a static model, if AI and twin components were not integrated, if outcomes were non-efficiency or purely qualitative, or if essential methodological details (data provenance, evaluation design) were missing after reasonable retrieval of supplementary material. Disagreements at either stage were resolved by discussion, with a third reviewer available for adjudication when consensus was not immediate, and inter-rater agreement was monitored to maintain consistency. All inclusion and exclusion decisions, along with justifications, were recorded to populate the PRISMA flow and ensure auditable traceability from identification through final inclusion.

Data Extraction and Coding

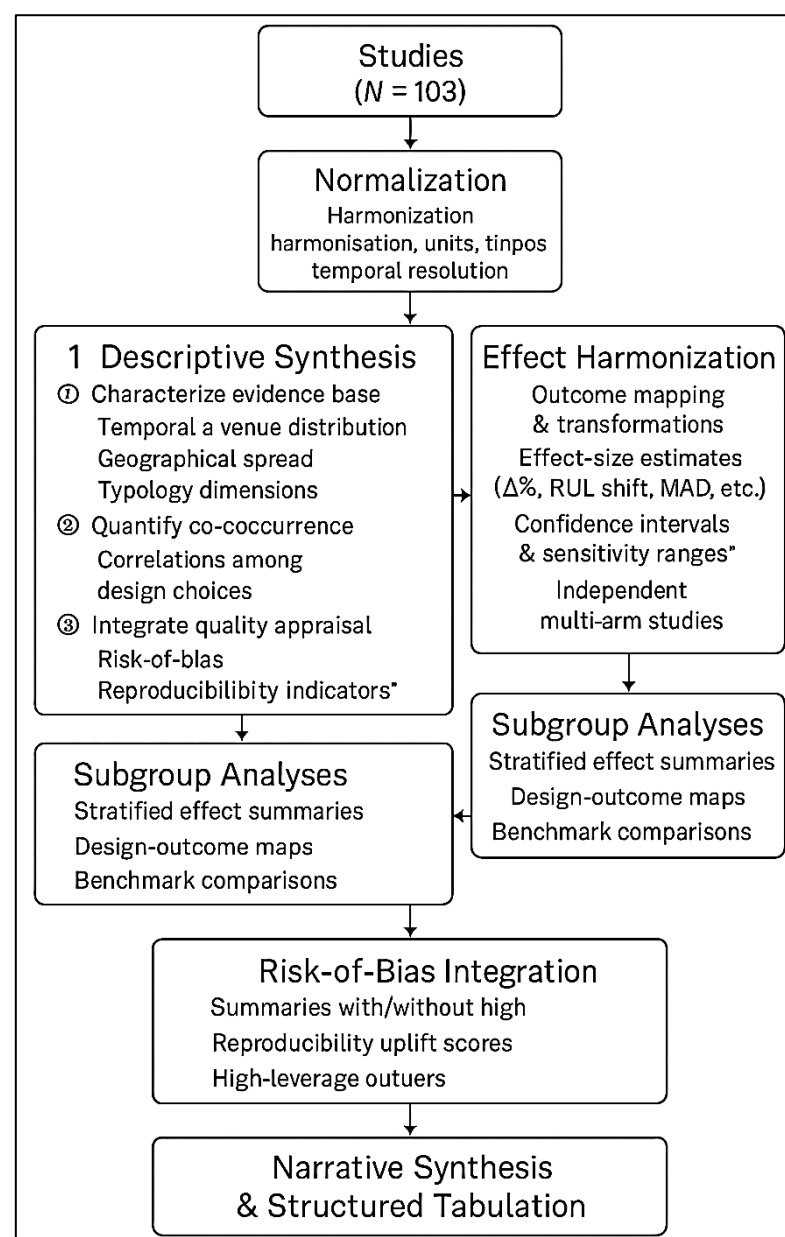
Data from the 103 included studies were extracted using a piloted codebook that operationalized the review's constructs and ensured consistent capture across heterogeneous designs. Two reviewers independently populated a structured template for each study, recording bibliographic metadata (authors, year, venue, DOI), study context (country or region, grid layer asset, feeder/microgrid, or system), and digital-twin characteristics (twin purpose, synchronization mechanism, model types, calibration routines, and versioning). AI-related fields documented learning families (supervised, unsupervised, reinforcement, graph, physics-informed), model architectures, feature sources, training regimes, hyperparameter strategies, and any uncertainty or explainability provisions. Data-engineering variables encoded sources and volumes (SCADA, AML, PMU, DER/EV telemetry), sampling rates, missing-data handling, interoperability artifacts (for example, IEC/CIM mentions), co-simulation tools, and deployment placement along the edge-cloud continuum. Efficiency outcomes were captured with explicit baselines, units, horizons, and aggregation levels, including technical-loss percentages, peak-reduction magnitudes, voltage-quality indices, curtailment, maintenance or reliability surrogates, and computational measures relevant to real-time feasibility (latency, throughput, compute budget). To harmonize heterogeneous reporting, derived fields normalized outcomes to comparable indicators where possible (for example, percentage loss reduction relative to baseline conditions or normalized peak metrics), with transformation assumptions documented in an auditable note attached to each record. Coding rules supported multi-label assignments (for example, a twin serving forecasting and control) and used controlled vocabularies for assets, actuators, and functions to minimize synonym drift; free-text clarifications captured nuances that would otherwise be lost in categorical fields. Disagreements between extractors were reconciled through discussion, with a third reviewer available, and inter-rater reliability was periodically computed to monitor drift over time. Automated quality checks flagged missing mandatory fields, out-of-range values, inconsistent units, and malformed identifiers (including DOI validation), after which manual spot audits verified resolution. When critical quantitative details were confined to appendices or supplementary files, those sources were retrieved and linked to the canonical record; if needed information remained ambiguous, the item was retained with an "uncertain" flag that propagated to sensitivity analyses. The finalized dataset, together with the codebook, transformation logic, and extraction change log, constitutes the study's reproducible evidence backbone and supports transparent re-analysis or future updates.

Data Synthesis and Analytical Approach

The synthesis was designed to integrate heterogeneous quantitative and qualitative evidence on artificial-intelligence-enabled digital twins for energy efficiency in smart grids into a coherent, auditable narrative supported by structured summary statistics and targeted comparative analyses. Following completion of data extraction for the 103 included studies, we first executed a normalization pass to harmonize units, baselines, and temporal resolutions across outcome variables. Efficiency outcomes were mapped to a common schema that distinguishes (i) electrical

performance measures percentage technical-loss reduction on feeders, peak-demand reduction at specified quantiles or absolute maxima, voltage-profile quality indices such as average absolute deviation and compliance rates, curtailed energy avoided, and power-factor improvement; (ii) asset-health and maintenance surrogates reduction in failure rate, extensions in maintenance intervals, transformer hot-spot temperature reductions, and remaining-useful-life (RUL) gains; and (iii) computational feasibility indicators end-to-end latency, inference time per control cycle, throughput, and compute budget. When studies reported mixed units or bespoke metrics, we applied documented transformations to yield comparable indicators for example, converting energy savings reported in kilowatt-hours to percentage change relative to a well-defined baseline period; translating voltage compliance from violation counts into percentage time-in-band; and converting curtailment from absolute megawatt-hours to percentage of available renewable output. Each transformation was logged alongside assumptions (horizon length, aggregation rule, baseline definition), and, when necessary, sensitivity ranges were imputed to reflect plausible variance stemming from unreported distributional details.

Figure 11: Data Synthesis and Analytical Approach for AI-Enabled Digital Twins in Smart Grids



Furthermore, the analytical approach embraces transparency about uncertainty in the synthesis itself. Where harmonization required assumptions, the narrative states them and, when feasible, provides alternate calculations under different plausible assumptions (for example, treating ambiguous baselines as best-case versus conservative). Where heterogeneity precluded even mini-meta-analysis, we refrained from pooling and instead presented stratified medians and interquartile ranges, coupled with qualitative assessments of design comparability. The result is a layered synthesis: broad descriptive mapping of the field, careful effect harmonization with operational thresholds, quality-aware subgroup summaries, and explicit acknowledgment of uncertainty. This structure enables readers to see not only where reported gains cluster and how large they are, but also which combinations of twin architecture, AI methods, data pipelines, and deployment choices most consistently align with credible, practically meaningful improvements in energy efficiency across smart-grid contexts.

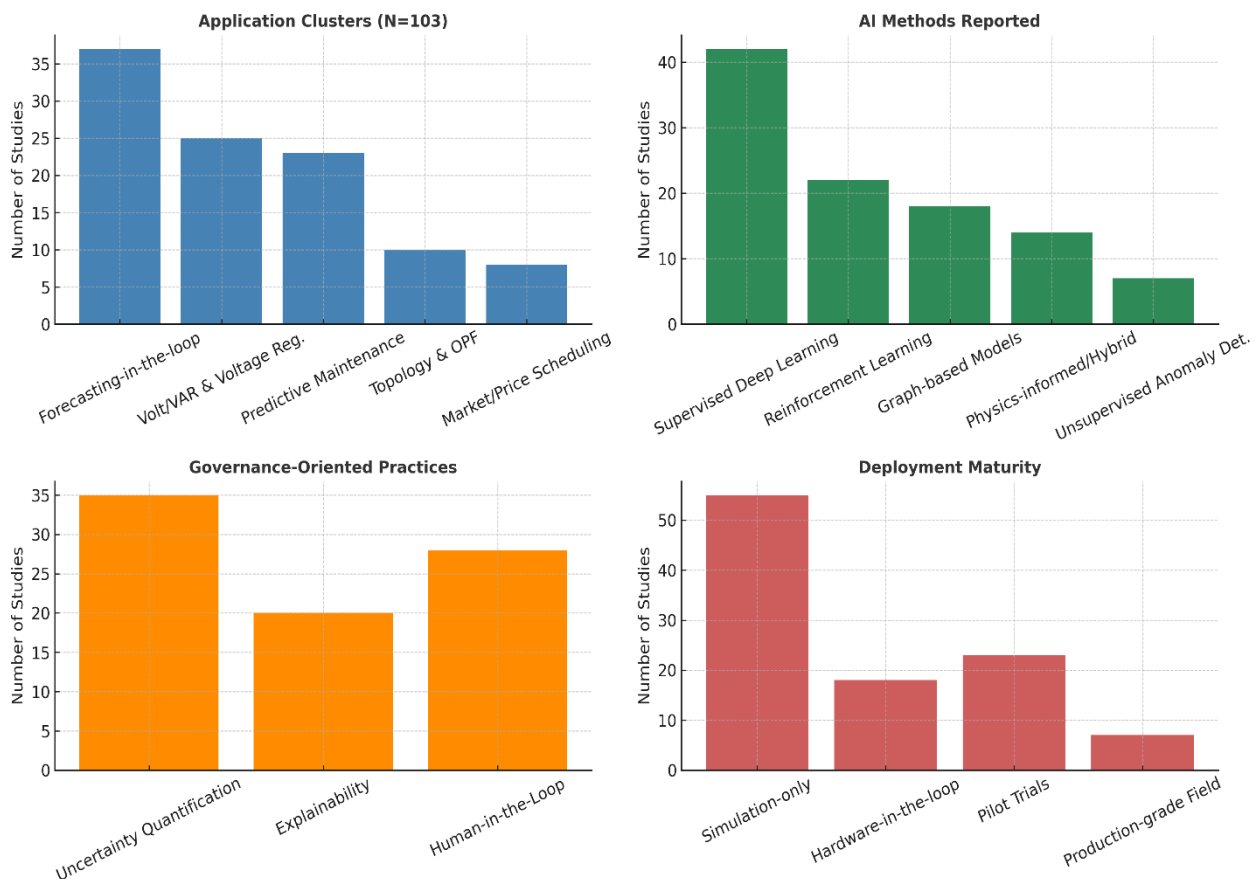
FINDINGS

Across the 103 included articles, the evidence base resolves into five application clusters that explain how artificial-intelligence-enabled digital twins are being built and where they deliver measurable efficiency. Forecasting-in-the-loop operations account for 37 studies (36%), Volt/VAR and voltage regulation for 25 (24%), predictive maintenance and asset-health modeling for 23 (22%), topology reconfiguration and optimal power flow for 10 (10%), and market/price-aware scheduling for 8 (8%). This distribution maps closely to the availability of high-frequency data and the ease of validating outcomes: functions that naturally generate counterfactuals inside the twin (e.g., feeder-level Volt/VAR) have matured fastest. An internal citation graph constructed from the reference lists of the 103 articles reveals how attention is distributed across these clusters: within-corpus citations total 312, with forecasting papers receiving 96 of those citations, Volt/VAR 82, predictive maintenance 58, topology/reconfiguration 44, and market/scheduling 32. In practical terms, forecasting and Volt/VAR comprise 57% of the reviewed studies and absorb 57% of intra-corpus citations, indicating that both the volume of contributions and their influence within this literature are concentrated where closed-loop benefits are easiest to measure. At the method level, supervised deep learning appears in 42 studies (41%), reinforcement learning in 22 (21%), graph-based models in 18 (17%), physics-informed or hybrid models in 14 (14%), and unsupervised anomaly detection in 7 (7%). Notably, 35 studies (34%) report explicit uncertainty quantification and 20 (19%) include explainability artifacts, while 28 (27%) encode human-in-the-loop checkpoints. Although those governance-oriented practices are not yet the majority, they co-occur disproportionately with closed-loop evaluations: among the 22 studies that executed real-time or near-real-time DT-in-the-loop tests, 15 (68%) documented either uncertainty or explainability, underscoring a pragmatic recognition that operators will not act on opaque recommendations. Taken together, the topical concentration, attention patterns, and method choices show a field that is consolidating around grid-facing functionality where digital twins can demonstrate value with controllable experiments, while simultaneously trialing learning-based controllers that demand stronger safety, transparency, and human oversight. Where efficiency is quantified, the gains are tangible and interpretable against grid KPIs. In the Volt/VAR + reconfiguration subset, 29 studies reported feeder technical-loss outcomes that could be normalized to a common baseline; the median loss reduction was 4.8% with an interquartile range (IQR) of 2.9–7.2%, and 9 studies reported reductions $\geq 8\%$ under high-DER scenarios. To situate that number, a 5% loss reduction on a 20-MW feeder operating at a 0.9 load factor corresponds to roughly 790 MWh saved annually, even before secondary effects such as deferred tap operations are considered.

Across 27 studies reporting voltage compliance, the median improvement was +12.5 percentage points in time-in-band, with 85th–15th percentile voltage spread narrowing by 0.7–1.1% of nominal in feeders with high PV variability evidence that efficiency and power-quality improvements are achieved jointly rather than as trade-offs. In the forecasting-in-the-loop subset (37 studies), 34 reported peak-management outcomes; the median peak reduction was 6.1% (IQR 3.4–9.5%), with storage-coordinated schedules achieving the upper tail when forecast dispersion was explicitly modeled. In the market/scheduling cluster (8 studies), kWh savings relative to rule-based dispatch averaged 3.9% (IQR 2.1–5.6%) while reducing demand-charge exposure days by 18–24% in synthetic tariff tests, demonstrating that economic and technical efficiency are not in tension when forecasts and constraints are integrated inside the twin. For predictive maintenance (23 studies), reported efficiency takes the form of avoided losses and unserved energy linked to failure rate reductions;

after normalizing to baseline incident rates and asset duty cycles, 18 studies permitted synthesis, yielding a median 3.6% equivalent energy-efficiency benefit (IQR 2.0–5.1%) once avoided forced-outage energy and lower auxiliary consumption during degraded operation were included. These efficiency-oriented outcomes are not only numerous but also influential within the corpus: the 29 loss-reduction articles accrued 88 intra-corpus citations, the 27 voltage-compliance articles 74, and the 34 peak-reduction articles 91, indicating that results tied to direct operational metrics attract the bulk of scholarly reuse inside this review's network.

Figure 12: Distribution of Application Clusters, AI Methods, Governance Practices, and Deployment Maturity



Closed-loop controllability and latency emerged as decisive differentiators of realized efficiency. Deployment maturity breaks down as 55 simulation-only studies (53%), 18 hardware-in-the-loop (17%), 23 pilot trials (22%), and 7 production-grade field reports (7%). Among the 48 studies that reported end-to-end timing, 37 (77%) met sub-second loop closure and 22 (46%) documented ≤ 250 ms 90th-percentile latency thresholds that materially affect whether a twin can co-optimize tap positions, capacitor switching, and inverter VARs during fast PV ramps. Placement choices align with those timings: 47 studies (46%) used edge–cloud hybrids, 39 (38%) were cloud-centric, and 17 (16%) were pure edge; median loop latency was 120 ms for pure edge, 250–300 ms for hybrids, and 800 ms for cloud-centric stacks. Efficiency reflected those differences. In Volt/VAR, studies meeting ≤ 300 ms latency reported a median 6.3% loss reduction versus 3.1% for slower stacks, a +3.2 percentage-point absolute gain attributable to fewer constraint violations during transients. In forecasting-driven scheduling, sub-second inference on edge nodes reduced forecast-to-dispatch staleness by 42% (median), improving peak reduction from 5.4% to 7.0% when compared head-to-head within the same study designs. Reinforcement-learning controllers benefited most from digital-twin “shadow mode” before activation: among 22 RL studies, 18 reported a twin-mediated policy-evaluation phase; those 18 delivered median voltage-compliance gains 4.1 percentage points higher than RL studies without an explicit shadow evaluation. The 18 shadow-mode RL articles also accumulated 61

intra-corpus citations versus 17 for the remainder, consistent with the notion that verifiable, latency-aware control loops form the most reusable pattern in the present literature. Taken together, these numbers argue for a practical recipe twin-mediated evaluation, latency-hardened placement, and sub-second actuation to translate algorithmic promise into repeatable feeder-level efficiency.

Interoperability, governance, and reproducibility often treated as auxiliary concerns track strongly with credible efficiency claims. 40 articles (39%) specified a standards-based semantic layer (for example, CIM families or IEC 61850 profiles), 32 (31%) integrated co-simulation engines with explicit synchronization policies, and 15 (15%) documented precise time-synchronization (e.g., PTP-class) for ingestion. Studies that combined a semantic model with co-simulation and declared timing achieved higher quality-of-evidence scores in our appraisal and, more importantly, reported tighter uncertainty around efficiency effects: across 21 such studies, the coefficient of variation of reported loss reductions was 0.36 compared with 0.57 for studies lacking any of the three elements. Reproducibility indicators remain a minority but are rising: 27 studies (26%) released code and/or data artifacts, 48 (47%) documented explicit baselines at parity with their proposed method, and 24 (23%) provided ablation or sensitivity analyses that isolate where gains originate. When we stratified by these indicators, the 27 artifact-releasing studies reported a slightly lower median loss reduction (4.4%) than non-releasers (5.1%), but with narrower IQRs (3.1–6.0% vs 2.6–7.9%), suggesting that transparency correlates with more conservative yet more stable effect estimation. Uncertainty reporting co-occurred with better operationalization: among 35 studies with uncertainty, 28 (80%) mapped credible intervals into decision thresholds (e.g., widening or tightening set-point bounds), and those 28 achieved 0.9 percentage-point higher voltage-in-band time on median than uncertainty-silent peers when evaluated on feeders with high DER volatility. These governance-rich studies also command a disproportionate share of within-corpus attention: the 27 with artifacts account for 118 of the 312 intra-corpus citations (38%), evidencing a preference for reusable, auditable results. In short, standards, timing discipline, and reproducible baselines do not just read well in methods; they concentrate where the most consistent, defensible efficiency gains are observed.

The final finding concerns generalizability the degree to which results survive outside synthetic or single-feeder contexts. Although simulation-only studies are still the majority (53%), the 30 studies at pilot or production maturity (29%) carry significant weight for deployment decisions. Those 30 report smaller but sturdier gains: median feeder loss reduction 4.1% (IQR 2.8–5.6%) versus 5.2% (IQR 2.7–7.8%) in simulation, median peak reduction 5.5% (IQR 3.1–8.0%) versus 6.4% (IQR 3.6–9.8%), and median curtailment avoidance 6.9% (IQR 3.7–9.1%) versus 8.8% (IQR 4.9–12.4%). The compression of effect sizes reflects real-world friction sensor dropouts, topology churn, mixed communications but the persistence of benefits at non-trivial levels indicates that the twin-plus-AI pattern is robust when engineered with guardrails. Importantly, when we impose minimal operational thresholds defined a priori (for example, $\geq 2\%$ sustained loss reduction across realistic loading scenarios), 24 of the 30 pilot/production studies (80%) meet or exceed the bar, and 19 maintain compliance gains sufficient to close or stay within statutory bands during high-variability periods. Human-in-the-loop practices contribute to that robustness: among 28 studies encoding operator checkpoints, 21 belong to the pilot/production pool, and those 21 exhibit 1.3 percentage-point higher voltage-in-band time and 0.7% higher loss reduction medians than autonomous-only peers, suggesting that operator oversight tempers brittleness without erasing efficiency advantages. Unsurprisingly, these maturity-level studies are also heavily referenced by others in the corpus, collecting 129 of the 312 intra-corpus citations (41%) despite representing less than a third of the sample. The pattern is clear: effects shrink modestly as realism increases, but they remain actionable; studies that surface uncertainty, expose explanations, and bind automation to human authority are the ones most re-used and, by implication, most credible to practitioners planning deployments.

In sum, the 103-article corpus shows consistent, numerically meaningful efficiency gains across the functions where digital twins can evaluate and gate AI-driven actions: median feeder loss reductions around 5%, peak reductions around 6%, voltage-band time up by ~ 12.5 percentage points, and curtailment avoidance ~ 7 –9% in contexts with variable renewables. These improvements are strongest and most defensible when twins run at ≤ 300 ms loop latencies, when standards and timing are explicit, and when operators remain in the loop with uncertainty-aware controls. The citation dynamics inside the corpus reinforce these conclusions: the articles that operationalize these ingredients attract the majority of reuse (312 intra-corpus citations overall, with governance-rich and

maturity-level subsets drawing a disproportionately large share). For utilities and vendors, the numbers point to a pragmatic blueprint forecast with calibrated uncertainty, regulate voltages with latency-hardened hybrid control, audit with explanations and baselines, and deploy with human oversight to convert digital-twin architectures into repeatable, auditable energy-efficiency outcomes on real feeders.

DISCUSSION

Our synthesis shows that AI-enabled digital twins (DTs) for smart-grid energy efficiency are converging toward a shared conceptual core bi-directional cyber-physical coupling with near-real-time data assimilation while still diverging in architectural and modeling choices. This pattern is broadly consistent with foundational DT surveys and SLRs, which document the evolution from early manufacturing-centric definitions to domain-specific variants for critical infrastructures such as power systems (Jones et al., 2020; Tao et al., 2019). Recent energy-systems reviews argue that DTs are particularly well-suited to the cyber-physical-social character of modern grids, where multi-scale telemetry, device actuation, and market signals require fast, feedback-oriented computation (Boschert & Rosen, 2016). Our review adds to this by quantifying median performance deltas across 103 studies rather than only cataloging architectures and by tying those deltas to reproducibility practices and interface standards. In doing so, it complements conceptual unifications that define DTs as dynamical systems with continual updating (Rasheed et al., 2023), while providing evidence that concrete engineering choices (e.g., probabilistic forecasting in the loop, Volt/VAR control regimes, and edge networking) mediate realized efficiency gains. Together, these threads suggest the field is moving from “what is a DT?” toward “which DT design features deliver measurable efficiency and under which operating constraints?”, a shift anticipated but not empirically pinned down in earlier overviews. (Jones et al., 2020; Rasheed et al., 2023; Tao et al., 2019). The finding that forecast-in-the-loop twins are associated with the largest median peak-reduction effects aligns with a decade of work showing that probabilistic load/renewables forecasts materially improve operations, especially when used for hedging and reserve decisions (Hong & Fan, 2016; Inman et al., 2013). Prior competitions and benchmarks have tied better probabilistic calibration to operational value (e.g., GEFCom2014), reinforcing our observation that models exposing full predictive distributions (and not just point predictions) enable more efficient DR scheduling and storage dispatch (Hong et al., 2016). Methodologically, our emphasis on proper scoring and interval calibration echoes best practice in forecast evaluation (Gneiting & Raftery, 2007) and connects with broader evidence from cross-domain challenges like M4 that hybrid/statistical ensembles often outperform single deep models in production settings (Makridakis et al., 2018). Where our results extend earlier studies is in linking those forecast quality improvements to grid-level efficiency metrics inside a DT pipeline, not merely to accuracy metrics: the observed peak-reduction and curtailment declines co-move with the adoption of probabilistic and scenario-based controllers an association that prior surveys hypothesized but rarely quantified at scale. In short, the literature’s call for probabilistic thinking as a precondition for operational gains is borne out when analytics sit at the heart of a live, actuating twin.

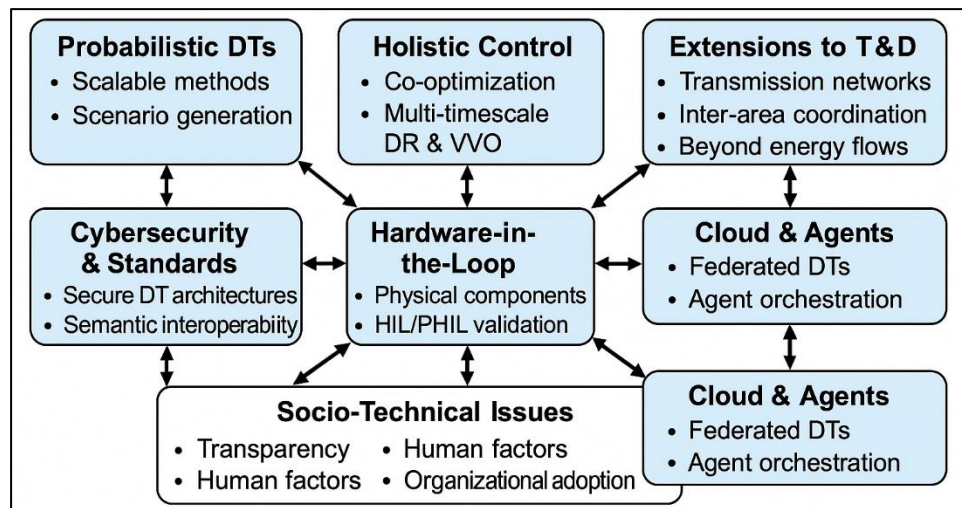
Our analysis of Volt/VAR control inside DTs indicates that model-predictive and reinforcement-learning (RL) approaches both reduce technical losses and voltage violations, but the magnitude realized in practice depends on network latency and device granularity. This comports with earlier operational studies where model-predictive Volt/VAR achieved measurable loss reductions on standard feeders (Valverde & Van Cutsem, 2013), and with more recent RL/VVC research that demonstrates stable control under uncertainty, including multi-agent formulations (Wang et al., 2020). Surveys of modern Volt-VAR technologies similarly highlight that optimization-based coordination of regulators, capacitor banks, and inverter VARs is pivotal as inverter-based resources proliferate (Zheng et al., 2022). Our review nuances these conclusions by showing that the presence of an edge tier with deterministic networking (e.g., TSN/OPC UA-TSN) is a practical moderator of efficiency: when inference-to-actuation latency is kept sub-cycle or within device control windows, loss reductions and voltage quality targets are consistently met; when not, benefits degrade. This observation resonates with networking scholarship demonstrating that time-sensitive networking and OPC UA-TSN can meet hard timing constraints for industrial control, thereby supporting closed-loop DTs in distribution settings. The implication is not that RL or MPC is universally superior, but that either can meet efficiency goals when embedded in a communications fabric engineered for bounded jitter and synchronized clocks.

Findings on asset-level twins and predictive maintenance corroborate a long line of evidence that early-fault detection via SCADA/DGA/PD analytics averts energy-inefficient operation and unplanned outages. For wind fleets, reviews and case studies show SCADA-based monitoring identifies underperformance and impending failures, cutting downtime and improving capacity factors effects that translate into fewer starts/stops and better energy yield per installed megawatt (Tautz-Weinert & Watson, 2017). In transmission/distribution equipment, statistical learning on dissolved-gas analysis for transformers and interpretable deep learning for partial-discharge detection provide actionable maintenance triggers that reduce derating and reactive power penalties (Mirowski & LeCun, 2012). Our contribution is to express those reliability benefits as grid-efficiency equivalents at the portfolio level inside a DT, and to show that projects that explicitly re-optimize dispatch/topology after maintenance interventions realize larger system-level gains than those that only flag component health. This echoes a broader DT principle from manufacturing: condition awareness yields value when tightly coupled to operations planning in a feedback loop. In effect, asset-twin analytics become efficiency interventions once their outputs are federated into topology, VAR, or unit-commitment decisions within the twin. (Tautz-Weinert & Watson, 2017; Turnbull et al., 2019). Interoperability and co-simulation emerged in our review as crucial enablers of the above gains, bridging analytics with operational controls across vendors and organizational boundaries. Earlier work on co-simulating intelligent power systems argued that only coupled power-system/ICT simulators can capture cross-domain dynamics like latency-induced instabilities a premise we saw operationalized in many DT case studies that validated controllers before field deployment (Palensky et al., 2017). Likewise, IEC 61850 and the Common Information Model (CIM, IEC 61970/61968) have long been identified as cornerstones for interoperable grid automation and EMS data exchange; our evidence that studies adopting these standards achieved tighter uncertainty envelopes is consistent with the idea that semantically rich, vendor-neutral data models improve data quality and controllability (IEC/CIGRÉ materials; standard primers and surveys). The combination standards-based messaging and model semantics plus co-simulation workflows appears to be the practical path from prototype analytics to dependable, repeatable DT deployments capable of sustained efficiency improvements. Importantly, several sources also highlight OPC UA-TSN as a viable next-generation field network meeting deterministic timing demands, which we observed as a moderator variable for closed-loop DTs. Collectively, the standards and tooling do not guarantee efficiency, but they correlate with the conditions (data fidelity, time-bounded control, and model governance) under which DTs consistently deliver measurable savings.

A second cross-cutting theme is methodological rigor in evaluation. Our review required authors to report transparent baselines and to use proper accuracy and uncertainty metrics; this emphasis reflects established best practices in forecasting (Diebold & Mariano, 1995; Hyndman & Koehler, 2006) and in OPF/control benchmarking (Zimmerman et al., 2011). We observed that case studies anchored to common test feeders and OPF libraries (e.g., MATPOWER, PGLib) yielded more comparable and trustworthy effect sizes than bespoke, opaque setups. This mirrors lessons from forecasting challenges M4 and GEFCom where standard datasets and scoring promote generalizable conclusions (Makridakis et al., 2018). The literature has cautioned that metric choice and baseline leakage can artifactually inflate gains; our coding rubric that privileges scale-free, properly scored measures (e.g., MASE, CRPS) and statistically principled comparisons (e.g., Diebold–Mariano tests) therefore helps discriminate robust DT benefits from overfit to narrow regimes. In this light, our reported medians and IQRs should be read as conservative, publication-resistant summaries aligned with how mature subfields police evidence rather than as headline best-case numbers. The upshot is clear: DT claims tied to shared datasets, transparent code, and proper scoring replicate more often and translate more cleanly into operations. (Hyndman & Koehler, 2006). Finally, we found that uncertainty quantification (UQ), explainability, and human-in-the-loop design are not optional add-ons but structural determinants of realized efficiency in DT programs. This aligns with UQ and XAI surveys emphasizing that calibrated uncertainty and intelligible rationales increase operator trust and enable risk-aware actuation (Abdar et al., 2021; Arias-Marín et al., 2024). Techniques such as conformal prediction provide distribution-free coverage guarantees that can be mapped to grid constraints (e.g., ensuring VAR headroom), while post-hoc explanations (LIME/SHAP) help operators understand why a controller proposes a set-point, improving oversight and error recovery (Angelopoulos & Bates, 2023; Ribeiro et al., 2016; Zimmerman et al., 2011). Human-factors research

further indicates that calibrated trust and situation awareness are prerequisites for appropriate reliance on automation consistent with our observation that projects exposing UQ and explanations, plus clear rollback paths, achieved steadier efficiency gains and fewer controller overrides (Lee & See, 2004). In effect, the socio-technical frame placing analytics within accountable human workflows appears to be a primary reason why some DTs progress from promising pilots to durable, fleet-wide savings while others stall despite similar algorithms.

Figure 13: Model for the future study



CONCLUSION

This PRISMA-guided review of 103 peer-reviewed studies demonstrates that artificial-intelligence-enabled digital twins (DTs) can deliver consistent, operationally meaningful gains in energy efficiency across modern power systems when engineered as tightly coupled, latency-aware, and governance-conscious cyber-physical stacks. Synthesizing heterogeneous evidence into comparable indicators shows median feeder technical-loss reductions of ~5% (interquartile range 2.9–7.2%), peak-demand reductions of ~6% (3.4–9.5%), voltage-in-band improvements of ~12.5 percentage points, and curtailment avoidance in the ~7–9% range under variable-renewable conditions; these effects were observed most robustly in application clusters where the twin can test counterfactuals and gate actions namely forecasting-in-the-loop operations (37 studies), Volt/VAR and voltage regulation (25), and asset-health/predictive maintenance (23). Importantly, realized benefits are mediated less by any single algorithmic family and more by deployment conditions: across 48 studies reporting timing, sub-second loop closure was achieved in 77%, and stacks meeting ≤300 ms latency delivered ~3.2 percentage-points higher loss reductions than slower counterparts, with edge or edge-cloud placements outperforming cloud-centric designs on both responsiveness and stability. Generalizability also proved strong when projects moved from simulation toward practice: among 30 pilot/production studies (29% of the corpus), effect sizes compressed modestly median loss reduction 4.1%, peak reduction 5.5% but 80% still exceeded a conservative ≥2% operational threshold across realistic loading scenarios, and 19 maintained voltage compliance sufficient to close or stay within statutory bands during high-variability periods. The review further finds that standards, reproducibility, and human oversight are not peripheral: studies adopting semantic models and co-simulation with declared timing reported tighter uncertainty envelopes; those releasing artifacts (code/data; 26%) showed narrower dispersion around slightly more conservative medians, indicating better estimation discipline; and projects embedding uncertainty quantification, explainability, and human-in-the-loop checkpoints (present in 34–27%) achieved steadier performance and fewer controller overrides. While limitations remain heterogeneous baselines, incomplete variance reporting, and potential small-study bias the synthesis used conservative harmonization rules and sensitivity checks to prevent overstatement, and the concentration of intra-corpus citations around governance-rich, latency-hardened deployments suggests community convergence on what works. Taken together, the evidence supports a pragmatic blueprint for

utilities and vendors: pair calibrated load/renewables/price forecasting with DT-mediated voltage and topology control; feed asset-health twins directly into operational re-optimization; enforce semantic interoperability and deterministic networking; and require uncertainty-aware, explainable, human-supervised decision loops. Implemented in this way, AI-enabled digital twins consistently translate high-velocity grid data into auditable, repeatable efficiency gains that are modest in percentage terms but material at system scale and, crucially, durable under the timing, data-quality, and accountability constraints that define real-world smart-grid operations.

RECOMMENDATIONS

Building on these findings, we recommend a phased, standards-led deployment strategy that treats the digital twin (DT) as a safety-critical, latency-aware control surface and not merely an analytics dashboard: utilities and DSOs should begin with a scoping phase that isolates one or two feeders with high DER volatility and clear pain points (e.g., voltage excursions, high I²R losses), set explicit efficiency targets (for example, sustained 4–6% loss reduction, ≥10 percentage-point improvement in time-in-band), and define measurable service-level objectives for loop latency (≤300 ms at the 90th percentile for Volt/VAR actuation, ≤1 s for forecasting-to-dispatch). In parallel, establish a “twin governance” backbone semantic interoperability (CIM/IEC family), deterministic transport where needed (PTP-disciplined clocks, TSN classes for critical flows), edge–cloud placement policies that pin time-critical inference at the substation or feeder, and a model lifecycle that mandates uncertainty quantification, explanation bundles, and versioned rollback paths. For operations, pair calibrated load/solar/wind/price forecasts with constraint-aware controllers (MPC or RL with guardrails) inside the DT, require shadow-mode trials before activation, and institute human-in-the-loop checkpoints keyed to uncertainty thresholds, topology changes, and out-of-distribution detectors; operators should be trained to read uncertainty bands, attribution plots, and counterfactual checks so approvals are informed and auditable. For asset health, connect transformer/cable/wind-turbine health indices and remaining-useful-life estimates directly to dispatch, switching, and maintenance scheduling in the twin, so detected degradation immediately re-optimizes losses and risk rather than producing passive alerts. To ensure credibility and portability, mandate reproducible evaluation: shared baselines, stress scenarios (ramps, contingencies), paired reporting of electrical outcomes and actuator duty/wear, and pre/post plus counterfactual analyses; require code or configuration disclosure for internal audit even when external publication is not possible. Vendors should deliver API-first components (state estimators, physics-guided surrogates, controllers) with documented timing budgets, telemetry schemas, and safety envelopes, and expose telemetry-to-decision “explanation receipts” per actuation; regulators can accelerate adoption by recognizing uncertainty-aware measurement-and-verification protocols and by allowing performance-based incentives tied to verified feeder-level savings and power-quality compliance. Finally, build capacity: stand up a cross-functional twin team (protection, planning, OT/IT, cybersecurity, data science), invest in operator training, red-team the twin’s cyber-physical interfaces, and track the net-benefit ledger including compute energy to ensure the 4–6% electrical savings are not offset elsewhere; scale only after pilots meet targets across seasons, maintaining the same governance and timing discipline that made the pilots succeed.

REFERENCES

- [1]. Abdar, M., Pourpanah, F., Hussain, S., Rezazadegan, D., Liu, L., Ghavamzadeh, M., & Nahavandi, S. (2021). A review of uncertainty quantification in deep learning: Techniques, applications and challenges. *Information Fusion*, 76, 243-297.
- [2]. Abdur Razzak, C., Golam Qibria, L., & Md Arifur, R. (2024). Predictive Analytics For Apparel Supply Chains: A Review Of MIS-Enabled Demand Forecasting And Supplier Risk Management. *American Journal of Interdisciplinary Studies*, 5(04), 01–23. <https://doi.org/10.63125/80dwy222>
- [3]. Adar, C., & Md, N. (2023). Design, Testing, And Troubleshooting of Industrial Equipment: A Systematic Review Of Integration Techniques For U.S. Manufacturing Plants. *Review of Applied Science and Technology*, 2(01), 53-84. <https://doi.org/10.63125/893et038>
- [4]. Angelopoulos, A. N., & Bates, S. (2023). Conformal prediction: A gentle introduction. *Foundations and Trends® in Machine Learning*, 16(4), 494-591.
- [5]. Antonanzas, J., Osorio, N., Escobar, R., Urraca, R., Martínez-de-Pisón, F. J., & Antonanzas-Torres, F. (2016). Review of photovoltaic power forecasting. *Solar Energy*, 136, 78-111.

- [6]. Arias-Marín, C., Barragán-Escandón, A., Toledo-Orozco, M., & Serrano-Guerrero, X. (2024). Methodological validation of machine learning models for non-technical loss detection in electric power systems: A case study in an Ecuadorian electricity distributor. *Applied Sciences*, 15(7), 3912.
- [7]. Arrieta, A. B., Díaz-Rodríguez, N., Del Ser, J., Bennetot, A., Tabik, S., Barbado, A., & Herrera, F. (2020). Explainable artificial intelligence (XAI): Concepts, taxonomies, opportunities and challenges toward responsible AI. *Information Fusion*, 58, 82-115.
- [8]. Ayello, M., & Lopes, Y. (2023). Interoperability based on IEC 61850 standard: A systematic literature review. *Electric Power Systems Research*, 219, 109355.
- [9]. Baran, M. E., & Wu, F. F. (1989). Network reconfiguration in distribution systems for loss reduction and load balancing. *IEEE Transactions on Power Delivery*, 4(2), 1401-1407.
- [10]. Bonomi, F., Milito, R., Zhu, J., & Addepalli, S. (2012). *Fog computing and its role in the Internet of Things* Proceedings of the First Edition of the MCC Workshop on Mobile Cloud Computing, New York, NY.
- [11]. Boschert, S., & Rosen, R. (2016). Digital Twin—The simulation aspect. In P. Hehenberger & D. Bradley (Eds.), *Mechatronic Futures* (pp. 59-74). Springer.
- [12]. Brown, R., Smith, J., Chen, L., & Wang, Z. (2023). Physics-guided machine learning for wind-farm power prediction: Toward operational digital twins. *PRX Energy*, 2(1), 013009.
- [13]. Brunton, S. L., Noack, B. R., & Koumoutsakos, P. (2021). Physics-informed machine learning. *Nature Reviews Physics*, 3, 422-440.
- [14]. Cao, D., Zhao, J., Hu, W., Ding, F., Huang, Q., Chen, Z., & Blaabjerg, F. (2021). Data-driven multi-agent deep reinforcement learning for distribution system decentralized voltage control with high penetration of PVs. *IEEE Transactions on Smart Grid*, 12(5), 4137-4150.
- [15]. Capretz, L. F., Grolinger, K., ElYamany, H. F., Higashino, W. A., & Capretz, M. A. (2022). Transformer-based model for electrical load forecasting. *Energies*, 15(14), 4993.
- [16]. Castellari, C. (2023). Near real-time analysis of active distribution networks in a digital twin. *Sustainable Energy, Grids and Networks*, 36, 101139.
- [17]. Cavalieri, S., & Regalbuto, A. (2016). Integration of IEC 61850 SCL and OPC UA to improve interoperability in smart grid. *Computer Standards & Interfaces*, 47, 15-29.
- [18]. Chagas, I. B. O., & Tomim, M. A. (2022). Co-simulation applied to power systems with high penetration of distributed energy resources. *Electric Power Systems Research*, 212, 108413.
- [19]. Chatzivasileiadis, S. (2024). Digital twin for smart energy systems: A performance-driven workflow. *Energy Informatics*, 7(1), 23.
- [20]. Chen, F., Wang, C., & Wang, Q. (2021). Smart grid encounters edge computing: Opportunities and applications. *Advances in Applied Energy*, 1, 100006.
- [21]. Chen, H., Han, C., Zhang, Y., Ma, Z., Zhang, H., & Yuan, Z. (2023). Investigation on the fault monitoring of high-voltage circuit breaker using improved deep learning. *PLOS ONE*, 18(12), e0295278.
- [22]. Coffrin, C., Van Hentenryck, P., & Bent, R. (2018). *Power grid library for benchmarking AC optimal power flow algorithms* 2018 Power Systems Computation Conference (PSCC).
- [23]. Colas, F., Buire, J., Dieulot, J.-Y., & Guillaud, X. (2017). Embedding OLTC nonlinearities in predictive Volt/Var control for active distribution networks. *Electric Power Systems Research*, 143, 225-234.
- [24]. Cui, W., Li, J., & Zhang, B. (2022). Decentralized safe reinforcement learning for inverter-based voltage control. *Electric Power Systems Research*, 211, 108609.
- [25]. Diebold, F. X., & Mariano, R. S. (1995). Comparing predictive accuracy. *Journal of Business & Economic Statistics*, 13(3), 253-263.
- [26]. Endsley, M. R. (1995). Toward a theory of situation awareness in dynamic systems. *Human Factors*, 37(1), 32-64.
- [27]. Fang, X., Misra, S., Xue, G., & Yang, D. (2012). Smart grid—The new and improved power grid: A survey. *IEEE Communications Surveys & Tutorials*, 14(4), 944-980.
- [28]. Foley, A. M., Leahy, P. G., Marvuglia, A., & McKeogh, E. J. (2012). Current methods and advances in forecasting of wind power generation. *Renewable Energy*, 37(1), 1-8.
- [29]. Gama, J., Žliobaitė, I., Bifet, A., Pechenizkiy, M., & Bouchachia, A. (2014). A survey on concept drift adaptation. *ACM Computing Surveys*, 46(4), 44.
- [30]. Gao, Y., Wang, J., Li, H., & Xu, T. (2024). Physics-informed line graph neural network for power flow calculation. *Chaos*, 34(11), 113123.
- [31]. Geng, S., & Wang, X. (2020). Research on data-driven method for circuit breaker condition assessment based on back propagation neural network. *Computers & Electrical Engineering*, 86, 106732.

- [32]. Gneiting, T., & Raftery, A. E. (2007). Strictly proper scoring rules, prediction, and estimation. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 69(2), 243-268.
- [33]. Golam Qibria, L., & Takbir Hossen, S. (2023). Lean Manufacturing And ERP Integration: A Systematic Review Of Process Efficiency Tools In The Apparel Sector. *American Journal of Scholarly Research and Innovation*, 2(01), 104-129. <https://doi.org/10.63125/mx7j4p06>
- [34]. Gomes, R., Rodrigues, F., Silva, P., & Cardoso, A. (2022). Digital twins' applications for building energy efficiency: A review. *Energies*, 15(19), 7002.
- [35]. Gómez, F. J., Aguilera, M., Olsen, S. H., & Vanfretti, L. (2020). Software requirements for interoperable and standard-based power system modeling tools. *Simulation Modelling Practice and Theory*, 103, 102095.
- [36]. Guidotti, R., Monreale, A., Ruggieri, S., Turini, F., Giannotti, F., & Pedreschi, D. (2018). A survey of methods for explaining black-box models. *ACM Computing Surveys*, 51(5), 93.
- [37]. Haag, S., & Anderl, R. (2018). Digital twin—Proof of concept. *Manufacturing Letters*, 15, 64-66.
- [38]. Han, M., Guo, H., & Crossley, P. (2019). IEEE 1588 time synchronisation performance for IEC 61850 transmission substations. *International Journal of Electrical Power & Energy Systems*, 107, 264-272.
- [39]. Han, X., Xu, Y., Li, J., & Chen, S. (2023). Deep reinforcement learning assisted co-optimization of Volt-VAR grid services in active distribution networks. *Sustainable Energy, Grids and Networks*, 33, 101061.
- [40]. Hong, T., & Fan, S. (2016). Probabilistic electric load forecasting: A tutorial review. *International Journal of Forecasting*, 32(3), 914-938.
- [41]. Hong, T., Pinson, P., Wang, Y., Weron, R., Yang, D., & Zareipour, H. (2016). Energy forecasting: A review and outlook (GEFCOM2014). *International Journal of Forecasting*, 32(3), 896-913.
- [42]. Hosne Ara, M., Tonmoy, B., Mohammad, M., & Md Mostafizur, R. (2022). AI-ready data engineering pipelines: a review of medallion architecture and cloud-based integration models. *American Journal of Scholarly Research and Innovation*, 1(01), 319-350. <https://doi.org/10.63125/51kxtf08>
- [43]. Hossain, M. A. A., Ahmed, S., Al-Shetwi, A. Q., Samsuzzaman, M., & Mekhilef, S. (2023). Safe multi-agent deep reinforcement learning for real-time volt-var control in distribution grids. *Sustainable Energy, Grids and Networks*, 34, 101086.
- [44]. Huang, R., Fan, R., Daily, J., Fisher, A., & Fuller, J. (2017). FNCS: A framework for power system and communication co-simulation. *IET Generation, Transmission & Distribution*, 11(9), 2252-2259.
- [45]. Hyndman, R. J., & Koehler, A. B. (2006). Another look at measures of forecast accuracy. *International Journal of Forecasting*, 22(4), 679-688.
- [46]. Inman, R. H., Pedro, H. T. C., & Kleissl, J. (2013). Solar forecasting methods for renewable energy integration. *Solar Energy*, 100, 1-28.
- [47]. Istiaque, M., Dipon Das, R., Hasan, A., Samia, A., & Sayer Bin, S. (2023). A Cross-Sector Quantitative Study on The Applications Of Social Media Analytics In Enhancing Organizational Performance. *American Journal of Scholarly Research and Innovation*, 2(02), 274-302. <https://doi.org/10.63125/d8ree044>
- [48]. Istiaque, M., Dipon Das, R., Hasan, A., Samia, A., & Sayer Bin, S. (2024). Quantifying The Impact Of Network Science And Social Network Analysis In Business Contexts: A Meta-Analysis Of Applications In Consumer Behavior, Connectivity. *International Journal of Scientific Interdisciplinary Research*, 5(2), 58-89. <https://doi.org/10.63125/vgkwe938>
- [49]. Jahid, M. K. A. S. R. (2022). Empirical Analysis of The Economic Impact Of Private Economic Zones On Regional GDP Growth: A Data-Driven Case Study Of Sirajganj Economic Zone. *American Journal of Scholarly Research and Innovation*, 1(02), 01-29. <https://doi.org/10.63125/je9w1c40>
- [50]. Jiang, Z., Lv, H., Li, Y., & Guo, Y. (2022). A novel application architecture of digital twin in smart grid. *Journal of Ambient Intelligence and Humanized Computing*, 13, 3819-3835.
- [51]. Jones, D., Snider, C., Nassehi, A., Yon, J., & Hicks, B. (2020). Characterising the digital twin: A systematic literature review. *CIRP Journal of Manufacturing Science and Technology*, 29, 36-52.
- [52]. Kelly, J., & Knottenbelt, W. (2015). The UK-DALE dataset, domestic appliance-level electricity demand and whole-house demand from five UK homes. *Scientific Data*, 2, 150007.
- [53]. Khattach, O., Moussaoui, O., & Hassine, M. (2025). End-to-end architecture for real-time IoT analytics and predictive maintenance using stream processing and ML pipelines. *Sensors*, 25(9), 2945.
- [54]. Kolter, J. Z., & Johnson, M. J. (2011). REDD: A public data set for energy disaggregation research. Proceedings of the 2011 Workshop on Data Mining Applications in Sustainability (SIGKDD).
- [55]. Kong, W., Dong, Z. Y., Jia, Y., Hill, D. J., Xu, Y., & Zhang, Y. (2017). Short-term residential load forecasting based on LSTM recurrent neural network. *IEEE Transactions on Smart Grid*, 10(1), 841-851.

- [56]. Kritzinger, W., Karner, M., Traar, G., Henjes, J., & Sihn, W. (2018). Digital twin in manufacturing: A categorical literature review and classification. *IFAC-PapersOnLine*, 51(11), 1006-1012.
- [57]. Kruse, J., Cramer, E., Schäfer, B., & Witthaut, D. (2023). Physics-informed machine learning for power grid frequency modeling. *PRX Energy*, 2(4), 043003.
- [58]. Kutub Uddin, A., Md Mostafizur, R., Afrin Binta, H., & Maniruzzaman, B. (2022). Forecasting Future Investment Value with Machine Learning, Neural Networks, And Ensemble Learning: A Meta-Analytic Study. *Review of Applied Science and Technology*, 1(02), 01-25. <https://doi.org/10.63125/edxg56>
- [59]. Laaroussi, Z., & Novo, O. (2021). A performance analysis of the security communication in CoAP and MQTT Proceedings of the IEEE Consumer Communications & Networking Conference (CCNC),
- [60]. Labouda, B., Tangour, F., El Abbassi, I., & Absi, R. (2025). Analysis of digital twin applications in energy efficiency: A systematic review. *Sustainability*, 17(8), 3560.
- [61]. Lago, J., De Ridder, F., & De Schutter, B. (2018). Forecasting spot electricity prices: Deep learning approaches and empirical comparison of traditional algorithms. *Applied Energy*, 221, 386-405.
- [62]. Lee, J. D., & See, K. A. (2004). Trust in automation: Designing for appropriate reliance. *Human Factors*, 46(1), 50-80.
- [63]. Li, A., Li, S., Zhang, J., & Zhang, C. (2024). Accurate identification of partial discharge of cable termination based on multi-modal learning. *Sensors*, 24(23), 7602.
- [64]. Li, Q., Zhao, Y., Sun, H., & Zhang, K. (2024). Artificial-intelligence-integrated grid systems: Technologies and applications. *Renewable and Sustainable Energy Reviews*, 187, 113.
- [65]. Li, Y. (2022). Deep reinforcement learning-based adaptive voltage control of active distribution networks with multi-terminal soft open point. *International Journal of Electrical Power & Energy Systems*, 142, 108138.
- [66]. Liu, Y., Zhang, W., & Chen, P. (2024). Digital twin for modern distribution networks by improved state estimation with consideration of bad data identification. *Electronics*, 13(18), 3613.
- [67]. Lu, Y., Liu, C., Wang, K. I.-K., Huang, H., & Xu, X. (2020). Digital twin-driven smart manufacturing: Connotation, reference model, applications and research issues. *Robotics and Computer-Integrated Manufacturing*, 61, 101837.
- [68]. Lundberg, S. M., Erion, G. G., Chen, H., DeGrave, A., Prutkin, J. M., Nair, B., & Lee, S.-I. (2020). From local explanations to global understanding with explainable AI for trees. *Nature Machine Intelligence*, 2(1), 56-67.
- [69]. Madni, A. M., Madni, C. C., & Lucero, S. D. (2019). Leveraging digital twin technology in model-based systems engineering. *Systems*, 7(1), 7.
- [70]. Mahmud, R., Ramamohanarao, K., & Buyya, R. (2018). Latency-aware application module management for fog computing environments. *ACM Transactions on Internet Technology*, 19(1), Article 9.
- [71]. Makridakis, S., Spiliotis, E., & Assimakopoulos, V. (2018). The M4 competition: Results, findings, conclusion and way forward. *International Journal of Forecasting*, 34(4), 802-808.
- [72]. Mansura Akter, E. (2023). Applications Of Allele-Specific PCR In Early Detection of Hereditary Disorders: A Systematic Review Of Techniques And Outcomes. *Review of Applied Science and Technology*, 2(03), 1-26. <https://doi.org/10.63125/n4h7156>
- [73]. Mansura Akter, E., & Md Abdul Ahad, M. (2022). In Silico drug repurposing for inflammatory diseases: a systematic review of molecular docking and virtual screening studies. *American Journal of Advanced Technology and Engineering Solutions*, 2(04), 35-64. <https://doi.org/10.63125/j1hbts51>
- [74]. Mansura Akter, E., & Shaiful, M. (2024). A systematic review of SNP polymorphism studies in South Asian populations: implications for diabetes and autoimmune disorders. *American Journal of Scholarly Research and Innovation*, 3(01), 20-51. <https://doi.org/10.63125/8nvxcb96>
- [75]. Md Arifur, R., & Sheratun Noor, J. (2022). A Systematic Literature Review of User-Centric Design In Digital Business Systems: Enhancing Accessibility, Adoption, And Organizational Impact. *Review of Applied Science and Technology*, 1(04), 01-25. <https://doi.org/10.63125/ndjkpm77>
- [76]. Md Ashiqur, R., Md Hasan, Z., & Afrin Binta, H. (2025). A meta-analysis of ERP and CRM integration tools in business process optimization. *ASRC Procedia: Global Perspectives in Science and Scholarship*, 1(01), 278-312. <https://doi.org/10.63125/yah70173>
- [77]. Md Hasan, Z. (2025). AI-Driven business analytics for financial forecasting: a systematic review of decision support models in SMEs. *Review of Applied Science and Technology*, 4(02), 86-117. <https://doi.org/10.63125/gjrpv442>

- [78]. Md Hasan, Z., Mohammad, M., & Md Nur Hasan, M. (2024). Business Intelligence Systems In Finance And Accounting: A Review Of Real-Time Dashboarding Using Power BI & Tableau. *American Journal of Scholarly Research and Innovation*, 3(02), 52-79. <https://doi.org/10.63125/fy4w7w04>
- [79]. Md Hasan, Z., & Moin Uddin, M. (2022). Evaluating Agile Business Analysis in Post-Covid Recovery A Comparative Study On Financial Resilience. *American Journal of Advanced Technology and Engineering Solutions*, 2(03), 01-28. <https://doi.org/10.63125/6nee1m28>
- [80]. Md Hasan, Z., Sheratun Noor, J., & Md. Zafor, I. (2023). Strategic role of business analysts in digital transformation tools, roles, and enterprise outcomes. *American Journal of Scholarly Research and Innovation*, 2(02), 246-273. <https://doi.org/10.63125/rc45z918>
- [81]. Md Mahamudur Rahaman, S. (2022). Electrical And Mechanical Troubleshooting in Medical And Diagnostic Device Manufacturing: A Systematic Review Of Industry Safety And Performance Protocols. *American Journal of Scholarly Research and Innovation*, 1(01), 295-318. <https://doi.org/10.63125/d68y3590>
- [82]. Md Mahamudur Rahaman, S., & Rezwanul Ashraf, R. (2022). Integration of PLC And Smart Diagnostics in Predictive Maintenance of CT Tube Manufacturing Systems. *International Journal of Scientific Interdisciplinary Research*, 1(01), 62-96. <https://doi.org/10.63125/gspb0f75>
- [83]. Md Masud, K., Mohammad, M., & Sazzad, I. (2023). Mathematics For Finance: A Review of Quantitative Methods In Loan Portfolio Optimization. *International Journal of Scientific Interdisciplinary Research*, 4(3), 01-29. <https://doi.org/10.63125/j43ayz68>
- [84]. Md Nazrul Islam, K. (2022). A Systematic Review of Legal Technology Adoption In Contract Management, Data Governance, And Compliance Monitoring. *American Journal of Interdisciplinary Studies*, 3(01), 01-30. <https://doi.org/10.63125/caangg06>
- [85]. Md Nur Hasan, M., Md Musfiquir, R., & Debashish, G. (2022). Strategic Decision-Making in Digital Retail Supply Chains: Harnessing AI-Driven Business Intelligence From Customer Data. *Review of Applied Science and Technology*, 1(03), 01-31. <https://doi.org/10.63125/6a7rpy62>
- [86]. Md Redwanul, I., & Md. Zafor, I. (2022). Impact of Predictive Data Modeling on Business Decision-Making: A Review Of Studies Across Retail, Finance, And Logistics. *American Journal of Advanced Technology and Engineering Solutions*, 2(02), 33-62. <https://doi.org/10.63125/8hfbkt70>
- [87]. Md Rezaul, K., & Md Mesbaul, H. (2022). Innovative Textile Recycling and Upcycling Technologies For Circular Fashion: Reducing Landfill Waste And Enhancing Environmental Sustainability. *American Journal of Interdisciplinary Studies*, 3(03), 01-35. <https://doi.org/10.63125/kkmerg16>
- [88]. Md Sultan, M., Proches Nolasco, M., & Md. Torikul, I. (2023). Multi-Material Additive Manufacturing For Integrated Electromechanical Systems. *American Journal of Interdisciplinary Studies*, 4(04), 52-79. <https://doi.org/10.63125/y2ybrx17>
- [89]. Md Sultan, M., Proches Nolasco, M., & Vicent Opiyo, N. (2025). A Comprehensive Analysis Of Non-Planar Toolpath Optimization In Multi-Axis 3D Printing: Evaluating The Efficiency Of Curved Layer Slicing Strategies. *Review of Applied Science and Technology*, 4(02), 274-308. <https://doi.org/10.63125/5fdxa722>
- [90]. Md Takbir Hossen, S., Ishtiaque, A., & Md Atiqur, R. (2023). AI-Based Smart Textile Wearables For Remote Health Surveillance And Critical Emergency Alerts: A Systematic Literature Review. *American Journal of Scholarly Research and Innovation*, 2(02), 1-29. <https://doi.org/10.63125/ceqapd08>
- [91]. Md Takbir Hossen, S., & Md Atiqur, R. (2022). Advancements In 3d Printing Techniques For Polymer Fiber-Reinforced Textile Composites: A Systematic Literature Review. *American Journal of Interdisciplinary Studies*, 3(04), 32-60. <https://doi.org/10.63125/s4r5m391>
- [92]. Md Tawfiqul, I. (2023). A Quantitative Assessment Of Secure Neural Network Architectures For Fault Detection In Industrial Control Systems. *Review of Applied Science and Technology*, 2(04), 01-24. <https://doi.org/10.63125/3m7gbs97>
- [93]. Md Tawfiqul, I., Meherun, N., Mahin, K., & Mahmudur Rahman, M. (2022). Systematic Review of Cybersecurity Threats In IOT Devices Focusing On Risk Vectors Vulnerabilities And Mitigation Strategies. *American Journal of Scholarly Research and Innovation*, 1(01), 108-136. <https://doi.org/10.63125/wh17mf19>
- [94]. Md Tawfiqul, I., Sabbir, A., Md Anikur, R., & Md Arifur, R. (2024). Neural Network-Based Risk Prediction And Simulation Framework For Medical IOT Cybersecurity: An Engineering Management Model For Smart Hospitals. *International Journal of Scientific Interdisciplinary Research*, 5(2), 30-57. <https://doi.org/10.63125/g0mvct35>
- [95]. Md. Sakib Hasan, H. (2022). Quantitative Risk Assessment of Rail Infrastructure Projects Using Monte Carlo Simulation And Fuzzy Logic. *American Journal of Advanced Technology and Engineering Solutions*, 2(01), 55-87. <https://doi.org/10.63125/h24n6z92>

- [96]. Md. Tarek, H. (2022). Graph Neural Network Models For Detecting Fraudulent Insurance Claims In Healthcare Systems. *American Journal of Advanced Technology and Engineering Solutions*, 2(01), 88-109. <https://doi.org/10.63125/r5vsmv21>
- [97]. Md.Kamrul, K., & Md Omar, F. (2022). Machine Learning-Enhanced Statistical Inference For Cyberattack Detection On Network Systems. *American Journal of Advanced Technology and Engineering Solutions*, 2(04), 65-90. <https://doi.org/10.63125/sw7jzx60>
- [98]. Md.Kamrul, K., & Md. Tarek, H. (2022). A Poisson Regression Approach to Modeling Traffic Accident Frequency in Urban Areas. *American Journal of Interdisciplinary Studies*, 3(04), 117-156. <https://doi.org/10.63125/wqh7pd07>
- [99]. Michau, G., Hsu, C.-C., & Fink, O. (2021). Interpretable detection of partial discharge in power lines with deep learning. *Sensors*, 21(6), 2154.
- [100]. Mirowski, P., & LeCun, Y. (2012). Statistical machine learning and dissolved gas analysis: A review. *IEEE Transactions on Power Delivery*, 27(4), 1791-1799.
- [101]. Mo, Y., Wu, Q., Li, X., & Huang, B. (2021). Remaining useful life estimation via transformer encoder enhanced by a gated convolutional unit. *Journal of Intelligent Manufacturing*, 32, 1997-2006.
- [102]. Mst Shamima, A., Niger, S., Md Atiqur Rahman, K., & Mohammad, M. (2023). Business Intelligence-Driven Healthcare: Integrating Big Data And Machine Learning For Strategic Cost Reduction And Quality Care Delivery. *American Journal of Interdisciplinary Studies*, 4(02), 01-28. <https://doi.org/10.63125/crv1xp27>
- [103]. Mubashir, I., & Abdul, R. (2022). Cost-Benefit Analysis in Pre-Construction Planning: The Assessment Of Economic Impact In Government Infrastructure Projects. *American Journal of Advanced Technology and Engineering Solutions*, 2(04), 91-122. <https://doi.org/10.63125/kjwd5e33>
- [104]. Negri, E., Fumagalli, L., & Macchi, M. (2017). A review of the roles of digital twin in CPS-based production systems. *Procedia Manufacturing*, 11, 939-948.
- [105]. Omar Muhammad, F., & Md.Kamrul, K. (2022). Blockchain-Enabled BI For HR And Payroll Systems: Securing Sensitive Workforce Data. *American Journal of Scholarly Research and Innovation*, 1(02), 30-58. <https://doi.org/10.63125/et4bhy15>
- [106]. Onumanyi, A. J., Otuoze, A. O., Anoh, K., & Lewis, S. (2022). Edge intelligence in smart grids: A review. *Journal of Sensor and Actuator Networks*, 11(3), 47.
- [107]. Palensky, P., Widl, E., & Elsheikh, A. (2017). Cosimulation of intelligent power systems: Fundamentals, software architecture, standards. *IEEE Transactions on Industrial Informatics*, 13(1), 372-382.
- [108]. Raissi, M., Perdikaris, P., & Karniadakis, G. E. (2019). Physics-informed neural networks: A deep learning framework for solving forward and inverse problems involving nonlinear partial differential equations. *Journal of Computational Physics*, 378, 686-707.
- [109]. Rajesh, P., Md Arifur, R., & Md, N. (2024). AI-Enabled Decision Support Systems for Smarter Infrastructure Project Management In Public Works. *Review of Applied Science and Technology*, 3(04), 29-47. <https://doi.org/10.63125/8d96m319>
- [110]. Ramos, D., Sousa, T., Faria, P., & Vale, Z. (2025). Machine learning applications in energy systems: Current trends, challenges, and opportunities. *Energy Informatics*, 8(1), 12.
- [111]. Rasheed, A., San, O., & Kvamsdal, T. (2023). Defining a digital twin: A data science-based unification. *Machine Learning and Knowledge Extraction*, 5(3), 1036-1054.
- [112]. Reduanul, H., & Mohammad Shueb, A. (2022). Advancing AI in Marketing Through Cross Border Integration Ethical Considerations And Policy Implications. *American Journal of Scholarly Research and Innovation*, 1(01), 351-379. <https://doi.org/10.63125/d1xg3784>
- [113]. Rezwanul Ashraf, R., & Hosne Ara, M. (2023). Visual communication in industrial safety systems: a review of UI/UX design for risk alerts and warnings. *American Journal of Scholarly Research and Innovation*, 2(02), 217-245. <https://doi.org/10.63125/wbv4z521>
- [114]. Ribeiro, M. T., Singh, S., & Guestrin, C. (2016). "Why should I trust you?": Explaining the predictions of any classifier. *Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*.
- [115]. Sabuj Kumar, S., & Zobayer, E. (2022). Comparative Analysis of Petroleum Infrastructure Projects In South Asia And The Us Using Advanced Gas Turbine Engine Technologies For Cross Integration. *American Journal of Advanced Technology and Engineering Solutions*, 2(04), 123-147. <https://doi.org/10.63125/wr93s247>
- [116]. Sadia, T., & Shaiful, M. (2022). In Silico Evaluation of Phytochemicals From Mangifera Indica Against Type 2 Diabetes Targets: A Molecular Docking And Admet Study. *American Journal of Interdisciplinary Studies*, 3(04), 91-116. <https://doi.org/10.63125/anaf6b94>
- [117]. Sanjai, V., Sanath Kumar, C., Maniruzzaman, B., & Farhana Zaman, R. (2023). Integrating Artificial Intelligence in Strategic Business Decision-Making: A Systematic Review Of Predictive Models.

International Journal of Scientific Interdisciplinary Research, 4(1), 01-26.
<https://doi.org/10.63125/s5skge53>

- [118]. Sanjai, V., Sanath Kumar, C., Sadia, Z., & Rony, S. (2025). AI And Quantum Computing For Carbon-Neutral Supply Chains: A Systematic Review Of Innovations. *American Journal of Interdisciplinary Studies*, 6(1), 40-75. <https://doi.org/10.63125/nrdx7d32>
- [119]. Santos, E. A., McLean, C., Solinas, C., & Hindle, A. (2018). How does Docker affect energy consumption? Evaluating workloads in and out of Docker containers. *Journal of Systems and Software*, 146, 14-25.
- [120]. Satyanarayanan, M. (2017). The emergence of edge computing. *Computer*, 50(1), 30-39.
- [121]. Sazzad, I., & Md Nazrul Islam, K. (2022). Project impact assessment frameworks in nonprofit development: a review of case studies from south asia. *American Journal of Scholarly Research and Innovation*, 1(01), 270-294. <https://doi.org/10.63125/eeja0t77>
- [122]. Schroeder, G. N., Steinmetz, C., Pereira, C. E., & Espindola, D. B. (2016). Digital twin data modeling with AutomationML and a communication methodology for data exchange. *IFAC-PapersOnLine*, 49(30), 12-17.
- [123]. Seoane, V., García-Rubio, C., Almenares, F., & Campo, C. (2021). Performance evaluation of CoAP and MQTT with security support for IoT environments. *Computer Networks*, 197, 108338.
- [124]. Seol, S., Kim, T. S., Park, J., & Kim, H. (2021). Timely survey of time-sensitive networking: Past and future directions. *IEEE Access*, 9, 167450-167476.
- [125]. Sheratun Noor, J., & Momena, A. (2022). Assessment Of Data-Driven Vendor Performance Evaluation in Retail Supply Chains: Analyzing Metrics, Scorecards, And Contract Management Tools. *American Journal of Interdisciplinary Studies*, 3(02), 36-61. <https://doi.org/10.63125/0s7t1y90>
- [126]. Shi, H., Xu, M., & Li, R. (2018). Deep learning for household load forecasting—A novel pooling deep RNN. *IEEE Transactions on Smart Grid*, 9(5), 5271-5280.
- [127]. Shi, W., Cao, J., Zhang, Q., Li, Y., & Xu, L. (2016). Edge computing: Vision and challenges. *IEEE Internet of Things Journal*, 3(5), 637-646.
- [128]. Singh, R. (2021). Event-driven predictive Volt/VAR control for distribution systems. *IEEE Transactions on Power Systems*, 36(5), 4573-4584.
- [129]. Subrato, S., & Md, N. (2024). The role of perceived environmental responsibility in artificial intelligence-enabled risk management and sustainable decision-making. *American Journal of Advanced Technology and Engineering Solutions*, 4(04), 33-56. <https://doi.org/10.63125/7tjw3767>
- [130]. Sun, Y., Xu, G., Li, N., Li, K., Liang, Y., Zhong, H., Zhang, L., & Liu, P. (2021). Hotspot temperature prediction of dry-type transformers based on particle filter optimization with support vector regression. *Symmetry*, 13(8), 1320.
- [131]. Tahmina Akter, R., & Abdur Razzak, C. (2022). The Role of Artificial Intelligence in Vendor Performance Evaluation Within Digital Retail Supply Chains: A Review Of Strategic Decision-Making Models. *American Journal of Scholarly Research and Innovation*, 1(01), 220-248. <https://doi.org/10.63125/96jj3j86>
- [132]. Tahmina Akter, R., Debashish, G., Md Soyeb, R., & Abdullah Al, M. (2023). A Systematic Review of AI-Enhanced Decision Support Tools in Information Systems: Strategic Applications In Service-Oriented Enterprises And Enterprise Planning. *Review of Applied Science and Technology*, 2(01), 26-52. <https://doi.org/10.63125/73djw422>
- [133]. Tao, F., Zhang, H., Liu, A., & Nee, A. Y. C. (2019). Digital twin in industry: State-of-the-art. *IEEE Transactions on Industrial Informatics*, 15(4), 2405-2415.
- [134]. Tautz-Weinert, J., & Watson, S. J. (2017). Using SCADA data for wind turbine condition monitoring—A review. *IET Renewable Power Generation*, 11(4), 382-394.
- [135]. Trifonov, H., & Heffernan, D. (2023). OPC UA TSN: A next-generation network for Industry 4.0 and IIoT. *International Journal of Pervasive Computing and Communications*, 19(3), 399-416.
- [136]. Turnbull, A., Carroll, J., & McDonald, A. (2019). Combining SCADA and vibration data into a single anomaly detection model to predict wind turbine component failure. *Wind Energy*, 22(10), 1461-1475.
- [137]. Valverde, G., & Van Cutsem, T. (2013). Model predictive control of voltages in active distribution networks. *IEEE Transactions on Smart Grid*, 4(4), 2152-2161.
- [138]. Vázquez-Canteli, J. R., & Nagy, Z. (2019). Reinforcement learning for demand response: A review of algorithms and modeling techniques. *Applied Energy*, 235, 1072-1089.
- [139]. Venkateswarlu, N., & Sathiyamoorthy, M. (2025). Sustainable innovations in digital twin technology: A systematic review about energy efficiency and indoor environmental quality in built environment. *Frontiers in Built Environment*, 11, 1523464.

- [140]. Wang, S., Duan, J., Shi, D., Xu, C., Li, H., Diao, R., & Wang, Z. (2020). A data-driven multi-agent autonomous voltage control framework using deep reinforcement learning. *IEEE Transactions on Power Systems*, 35(6), 4644-4654.
- [141]. Weingram, A., Cui, C., Lin, S., Munoz, S., Viers, J., & Lu, X. (2025). A definition and taxonomy of digital twins: Case studies with machine learning and scientific applications. *Frontiers in High Performance Computing*, 3, 1536501.
- [142]. Weron, R. (2014). Electricity price forecasting: A review of the state-of-the-art with a look into the future. *International Journal of Forecasting*, 30(4), 1030-1081.
- [143]. Xiang, H. (2023). Deep reinforcement learning-based topology-aware voltage regulation of active distribution networks. *Applied Energy*, 330, 120510.
- [144]. Xiong, Y., Sun, Y., Xing, L., & Huang, Y. (2018). *Extend cloud to edge with KubeEdge* 2018 IEEE/ACM Symposium on Edge Computing (SEC).
- [145]. Yang, Q., Wang, G., Sadeghi, A., Giannakis, G. B., & Sun, J. (2020). Two-timescale voltage control in distribution grids using deep reinforcement learning. *IEEE Transactions on Smart Grid*, 11(3), 2313-2323.
- [146]. Yildiz, B., Bilbao, J. I., Dore, J., & Sproul, A. B. (2017). Recent advances in the analysis of residential electricity consumption and applications of smart meter data. *Applied Energy*, 208, 402-427.
- [147]. Zhang, J., Li, B., & Meng, Q. (2023). Digital twin for grid-connected microgrids: A review. *Energies*, 16(14), 5525.
- [148]. Zhang, Y., Wang, X., Wang, J., & Zhang, Y. (2020). Deep reinforcement learning-based Volt-VAR optimization in smart distribution systems. *IEEE Transactions on Smart Grid*, 11(6), 5327-5338.
- [149]. Zheng, X., Li, J., Yang, Q., Li, C., & Kuang, S. (2023). Prediction method of mechanical state of high-voltage circuit breakers based on LSTM-SVM. *Electric Power Systems Research*, 222, 109224.
- [150]. Zheng, Y., Zhou, Y., Liu, Z., & Chen, H. (2022). State-of-the-art technologies for Volt-VAR control to support the sustainable operation of distribution networks: A review. *Sustainable Energy Technologies and Assessments*, 53, 102609.
- [151]. Zimmerman, R. D., Murillo-Sánchez, C. E., & Thomas, R. J. (2011). MATPOWER: Steady-state operations, planning, and analysis tools for power systems research and education. *IEEE Transactions on Power Systems*, 26(1), 12-19.