



GPU-Accelerated Physics-Informed Digital Twins for Real-Time State Estimation and Fault Localization in Distribution Grids

Zaheda Khatun¹; Md Hamidur Rahman²;

[1]. Master of Engineering in Electrical and Computer Engineering(Continuing), Lamar University, USA;
Email: zahedadisha@gmail.com

[2]. Master of Engineering in Electrical and Computer Engineering(Continuing), Lamar University, USA;
Email: hamidurr99@gmail.com

Doi: [10.63125/msrpfb04](https://doi.org/10.63125/msrpfb04)

Received: 23 September 2024; **Revised:** 27 October 2024; **Accepted:** 29 November 2024; **Published:** 28 December 2024

Abstract

This study addresses a persistent operational problem in distribution-grid management: utilities need fast, trustworthy state estimation and fault localization to support restoration and switching decisions, yet conventional analytics are often too slow for real-time use and can be difficult to trust under changing grid conditions. The purpose of the study was to evaluate whether GPU-accelerated, physics-informed digital twins can measurably improve real-time state estimation, fault localization effectiveness, and operational decision support in cloud and enterprise deployment contexts. Using a quantitative, cross-sectional, case-based design, data were collected from a sample of $N = 210$ respondents drawn from cloud and enterprise distribution-grid digital-twin cases. The key independent variables were GPU Acceleration Capability, Physics-Informed Modeling Strength, Digital Twin Fidelity, and Data Integration Quality; the dependent variables were Real-Time State Estimation Performance, Fault Localization Effectiveness, and Operational Decision Support Value, with an additional trust indicator for physics-consistency confidence. The analysis plan employed reliability testing (Cronbach's alpha), Pearson correlations, and multiple regression modeling to estimate direct effects and the enabling pathway from acceleration and physics constraints to operational value. Reliability results indicated strong internal consistency across constructs ($\alpha = .83$ to $.91$). Correlation results showed that GPU capability was positively associated with state estimation ($r = .56, p < .001$) and that physics-informed strength had an even stronger association with state estimation ($r = .61, p < .001$); state estimation was strongly related to fault localization ($r = .63, p < .001$). Regression findings confirmed that state estimation performance was significantly predicted by GPU capability ($\beta = .29, p < .001$) and physics-informed strength ($\beta = .37, p < .001$), explaining 54% of variance ($R^2 = .54$). Fault localization effectiveness was significantly predicted by state estimation ($\beta = .41, p < .001$), digital twin fidelity ($\beta = .28, p < .001$), and data integration quality ($\beta = .19, p = .002$), explaining 62% of variance ($R^2 = .62$). Operational decision support value was driven primarily by fault localization ($\beta = .46, p < .001$) and state estimation ($\beta = .21, p = .004$), with GPU providing incremental contribution ($\beta = .14, p = .030$), explaining 58% of variance ($R^2 = .58$). Trust results further indicated high physics-consistency confidence ($M = 4.12, SD = 0.48$). Overall, the study implies that utilities should prioritize physics-consistent modeling and GPU-ready architectures, while investing in data integration quality and twin fidelity to convert computational speed into reliable, actionable grid decisions.

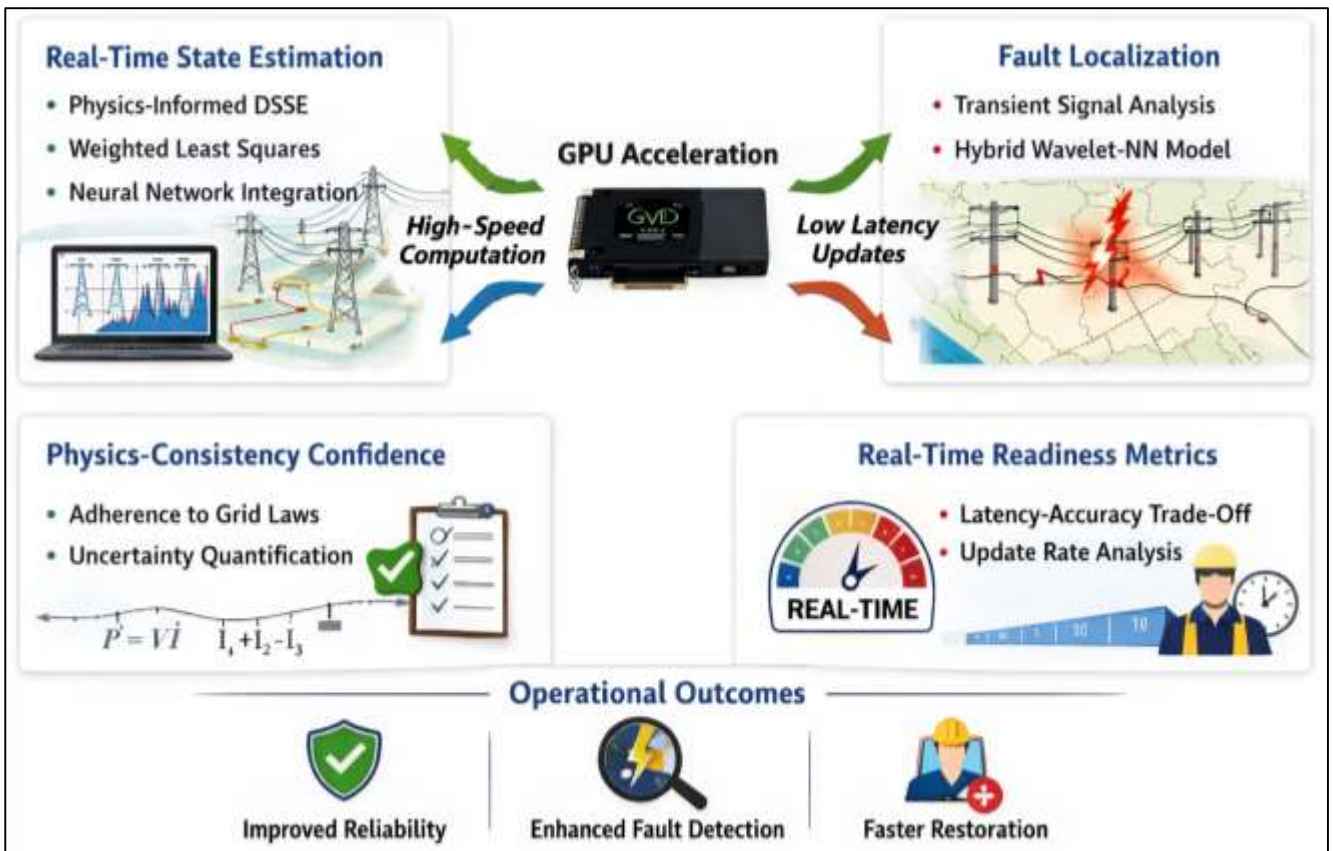
KEYWORDS

GPU Acceleration; Physics-Informed Digital Twins; Real-Time State Estimation; Fault Localization; Operational Decision Support;

INTRODUCTION

A digital twin is commonly defined as a continuously updated, data-linked virtual representation of a physical asset or process, designed to mirror system behavior through a combination of models, measurements, and computational orchestration (Adewole et al., 2016). In electric power engineering, the “asset” is often not a single device but an interconnected distribution grid, where network physics, operational constraints, and measurement uncertainty interact across space and time (Borghetti et al., 2008). A distribution grid refers to the medium- and low-voltage portion of the power system delivering electricity from substations to end users, typically characterized by radial or weakly meshed topology, phase imbalance, heterogeneous conductor types, and fast-changing net loads due to distributed energy resources and electrification (Borghetti et al., 2006).

Figure 1: GPU-Accelerated Physics-Informed Digital Twin Framework



The operational state of such a grid is commonly expressed through bus voltage magnitudes and phase angles, branch flows, and injection quantities. State estimation is the computational process of inferring these internal state variables from imperfect measurements; in practice, it often relies on variants of weighted least squares, constrained optimization, and hybrid measurement fusion (Pan & Liu, 2020). The distribution-level problem is notably harder than transmission-level estimation because distribution measurements are sparser, topology can be uncertain, loads are less predictable, and phase couplings are more pronounced. In parallel, fault localization refers to identifying the most likely faulted segment and estimating its location along a feeder, using signals such as voltage/current transients, steady-state phasors, or derived features; it is central to reliability restoration, safety, and outage management (Haque & Arifur, 2020; Raissi et al., 2019; Rauf, 2018). Signal-processing approaches have shown that fault-generated transients encode path-dependent signatures that can help localize faults. Meanwhile, data-driven classification and section identification strategies that fuse engineered features with learning models have demonstrated strong performance in complex feeder topologies. These foundational definitions frame why the proposed topic—GPU-accelerated physics-informed digital twins for real-time state estimation and fault localization—sits at the intersection of cyber-physical modeling, high-performance computing, and distribution automation, with global

significance because distribution systems carry the last-mile reliability burden for households, healthcare, manufacturing, and critical infrastructure (Zhang et al., 2019).

At an international scale, distribution grids are increasingly expected to support high penetrations of variable distributed generation, electric vehicles, advanced power electronics, and active demand response, all while maintaining tight reliability targets and operational resilience. This pushes utilities toward decision systems that can operate at operational timescales measured in seconds rather than minutes, motivating digital-twin architectures that can ingest streaming telemetry and run analytics continuously (Haque & Arifur, 2021; Ashraful et al., 2020; Zamzam & Sidiropoulos, 2020). The global challenge is not merely “more data,” but heterogeneous and imperfect data: supervisory control and data acquisition (SCADA) remains limited in many distribution contexts, advanced metering infrastructure (AMI) yields massive but delayed and noisy measurements, and newer sensors (phasor measurement units, micro-PMUs, power quality monitors) produce high-frequency signals that are rich but not universally deployed (Jinnat & Kamrul, 2021; Li & Li, 2014; Fokhrul et al., 2021). Consequently, the reliability of any digital twin depends on the fidelity of (i) the underlying physics model, (ii) the measurement-to-state mapping, and (iii) computational latency and robustness under non-ideal data. Literature on digital twin technology emphasizes that synchronization between the physical and digital counterparts requires networking and compute strategies that preserve real-time responsiveness and data quality. In power-system contexts, early digital-twin discussions highlight the need for “closed-loop” coupling between grid behavior and virtual analytics so that monitoring and control functions can operate on consistent representations (Hammad, 2022; Mashaly, 2021; Zaman et al., 2021). These requirements become operationally strict in distribution grids because faults, switching actions, and inverter dynamics can evolve quickly; thus, a digital twin that is not computationally timely can become informationally stale. In this landscape, the proposed research title directly addresses a central international pain point: achieving real-time state estimation and fault localization without sacrificing physics consistency, and doing so in a way that is measurable, testable, and scalable using GPU acceleration and physics-informed learning. The international relevance is also methodological: modern utilities and regulators increasingly demand evidence-based validation and auditable performance claims, which aligns with a quantitative, cross-sectional, case-study design that can statistically test relationships among readiness, physics-consistency, latency, and diagnostic accuracy (Mora-Flórez et al., 2008).

Distribution system state estimation (DSSE) has matured through decades of optimization-based approaches, yet its distribution-specific challenges persist. Traditional state estimation generally treats the grid as a set of nonlinear equations that relate measurements to unknown states, solved iteratively under a measurement error model (Jabed Hasan & Waladur, 2022; Arifur & Haque, 2022; Zhou et al., 2021). However, distribution networks complicate this process because measurement redundancy is weaker and the system is less observable without pseudo-measurements or strong priors. A widely cited direction is the integration of learning-based components that reduce online computational burden by shifting effort to offline training, provided that the learned mapping remains physically meaningful and robust. Physics-aware or physics-informed machine learning aims to preserve the structural constraints of the grid—such as power flow relations—while exploiting data to improve speed and resilience to poor initialization (Towhidul et al., 2022; Rifat & Jinnat, 2022). One influential example is the use of physics-aware neural networks for DSSE, where learning is guided by grid constraints to improve stability and computational efficiency relative to purely numerical solvers. In parallel, physics-informed neural networks (PINNs) were introduced as a general framework for embedding governing equations into neural training objectives, originally formulated to solve forward and inverse problems under nonlinear PDE constraints (Abdulla & Majumder, 2023; Borghetti et al., 2006; Rifat & Alam, 2022). While PINNs were not invented for power systems, their logic—penalizing violations of physics during learning—maps naturally to grid estimation problems where Kirchhoff’s laws and power flow relationships are the “physics.” Subsequent work on uncertainty in PINNs further clarified that physics-constrained learning must still account for data noise, model mismatch, and epistemic uncertainty if predictions are to be operationally trustworthy. Taken together, these strands motivate a digital-twin approach where DSSE is not only fast but also physics-consistent, because the digital twin is intended to be a faithful operational mirror. In this thesis context, physics-informed DSSE

becomes more than an algorithmic choice: it becomes a credibility mechanism, enabling the results section to report explicit evidence of physics-consistency confidence rather than relying solely on accuracy metrics (Mashaly, 2021).

Fault localization is the second pillar of operational intelligence in distribution grids and is tightly coupled to state estimation quality. A fault's location determines switching plans, restoration sequencing, crew dispatch, safety controls, and customer outage duration (Raissi et al., 2019). Classical fault-location techniques include impedance-based estimators and traveling-wave methods, each with known sensitivities to feeder laterals, load uncertainty, and measurement placement. Wavelet-based transient analysis demonstrated that fault-originated traveling waves generate frequency-time signatures that can be correlated with paths in the network, enabling faulted-section inference in complex medium-voltage networks (Faysal & Bhuya, 2023; Habibullah & Aditya, 2023; Zhou et al., 2021). Related wavelet-driven fault-location work also appears in power-system venues, reflecting continued interest in exploiting transients for localization. Meanwhile, hybrid learning approaches have been developed to address the "multiple estimation" ambiguity in branched feeders, where several laterals can correspond to similar apparent distances. For example, a two-stage method combining discrete wavelet transform features and neural networks was proposed for fault section identification and fault location, explicitly targeting distribution topology complexity (Hammad & Mohiul, 2023; Haque & Arifur, 2023). Comparative studies of fault location methods emphasize that distribution networks demand techniques resilient to uncertainty in load, fault resistance, and sparse measurement coverage. In the conceptual frame of a physics-informed digital twin, fault localization can be treated as both an inference problem (pattern recognition from signals) and a model-consistency problem (ensuring that localized fault hypotheses align with feasible grid states) (Akbar & Farzana, 2023; Mostafa, 2023). That duality is important because a digital twin that localizes faults without reconciling them against state-estimation physics may produce operationally brittle decisions (Raissi et al., 2019; Zamzam & Sidiropoulos, 2020). Therefore, the proposed thesis focus implicitly connects DSSE and fault localization as mutually reinforcing: improved state estimates provide cleaner residuals and consistency checks for fault hypotheses, while fault-aware modeling improves state estimation by explaining abrupt measurement deviations (Borghetti et al., 2006).

A key barrier to real-time digital twins is computational latency under large-scale nonlinear inference and repeated linear algebra. GPU acceleration is widely recognized as a practical route to improving throughput for matrix-heavy workloads, especially in iterative solvers and sparse computations common in power systems. For example, GPU-based approaches have been applied to accelerate power flow-related computations using polynomial preconditioning and conjugate-gradient methods, demonstrating notable performance potential for large systems (Jahangir & Hammad, 2024; Rifat & Rebeka, 2023). More directly tied to state estimation, a GPU-based matrix-structure-driven strategy for weighted least squares state estimation was reported to significantly speed computation by separating structural preprocessing (CPU) from repeated numerical kernels (GPU), leveraging sparse patterns and tuned parallel operations (Li & Li, 2014; Masud & Hammad, 2024; Md & Praveen, 2024). The conceptual importance of such work for this thesis is not limited to "faster compute," but to enabling a real-time operational loop where the digital twin can update frequently enough to remain synchronized with the physical grid (Raissi et al., 2019; Zhang et al., 2019). If a digital twin is defined as a living model that is updated continuously through data and computation, then latency is not a secondary performance metric; it is a definitional requirement (Rifat & Rebeka, 2024; Sai Praveen, 2024). Moreover, latency and accuracy may trade off: aggressive approximation can increase speed but degrade estimation fidelity, while high-fidelity nonlinear solvers can be too slow for real-time use. This is precisely why the results sections you proposed – such as a Latency–Accuracy Trade-Off Evidence block – fit the study uniquely: they operationalize "real-time readiness" as a measurable construct rather than an informal claim. In addition, GPU acceleration interacts naturally with physics-informed learning: once trained, neural inference can be extremely fast, and GPUs can support both training (offline) and inference (online) (Shehwar & Nizamani, 2024; Azam & Amin, 2024). The thesis theme therefore becomes a coherent pipeline: physics-informed models to maintain credibility, GPU acceleration to maintain real-time synchronization, and quantitative evidence to demonstrate that the combined system achieves dependable operational performance in a realistic case-study context (Borghetti et al., 2006).

Trustworthiness for a physics-informed digital twin is not only computational; it is epistemic. Distribution grids are uncertain systems: parameter errors, topology errors, unmodeled control actions, meter bias, and missing data can all cause model–reality divergence. Physics-informed learning addresses part of the problem by penalizing violations of governing relations, but uncertainty quantification remains central if a digital twin is expected to support decisions under risk (Pan & Liu, 2020). Research on uncertainty in physics-informed neural networks has explicitly argued that forward and inverse predictions must incorporate uncertainty sources to avoid overconfident outputs when data are sparse or noisy (Zhang et al., 2019). This aligns directly with state estimation practice, where measurement error covariances and bad-data detection are foundational. In addition, the DSSE context raises domain-specific concerns: unbalanced three-phase modeling, pseudo-measurement construction, and correlated noise in AMI-derived features can complicate classical assumptions. Thus, physics-consistency metrics can be interpreted as a type of “model credibility monitor,” complementing statistical accuracy (Pan & Liu, 2020). In your outline, the Physics-Consistency Confidence Results section is therefore not decorative; it represents a study-specific credibility argument: the digital twin’s outputs are not only accurate against benchmarks but also consistent with the physical constraints that define feasible grid states. Similarly, the Real-Time Readiness Index can be grounded in measurable latency percentiles, update rates, and convergence reliability across scenarios, tying directly to the digital twin’s requirement for synchronous operation (Mora-Flórez et al., 2008). These study-specific constructs also support your quantitative design: they can be operationalized as Likert-scale constructs captured from engineers/operators (perceived readiness, interpretability, actionability) and paired with objective system metrics (runtime, residual norms, constraint violation), enabling correlation and regression analysis to test hypotheses about how GPU acceleration and physics-informed design choices relate to trust, usability, and performance outcomes (Li & Li, 2014).

Within this conceptual foundation, the motivation for the proposed thesis title can be positioned as a structured response to three coupled demands: real-time operation, physics-grounded credibility, and diagnostic usefulness (state estimation plus fault localization) in distribution grids. Digital twin research highlights the importance of networking, synchronization, and low-latency data pathways for keeping virtual models aligned with physical assets. Power-system digital twin discussions in the IEEE conference ecosystem similarly emphasize platform-level design and dispatching applications, reflecting the operational orientation of the field (Borghetti et al., 2008). On the algorithmic side, physics-aware learning for DSSE demonstrates that embedding physical structure can stabilize estimation and improve real-time suitability compared to purely iterative nonlinear solvers (Pan & Liu, 2020). Meanwhile, foundational PINN work provides a widely recognized methodological template for enforcing governing relations in learning systems, and uncertainty-focused extensions clarify that confidence must be quantified rather than assumed. On the fault localization side, traveling-wave and wavelet-based methods demonstrate that transient signals encode locational information, while hybrid wavelet–neural approaches explicitly target topology-driven ambiguity in branched feeders. Finally, GPU-oriented work shows that the computational bottlenecks of iterative estimation and related matrix operations can be reduced substantially with structure-aware GPU kernels and heterogeneous scheduling (Raissi et al., 2019). These strands collectively justify a research design that tests, in a concrete case-study environment, whether a GPU-accelerated physics-informed digital twin can deliver statistically defensible improvements in readiness, physics consistency, and diagnostic performance for DSSE and fault localization (Mora-Flórez et al., 2008).

This study is designed to achieve a set of clearly defined objectives that align directly with the research title, the quantitative cross-sectional strategy, and the case-study grounding in distribution-grid operations. The first objective is to operationalize the core capabilities of a GPU-accelerated physics-informed digital twin into measurable constructs that can be examined statistically, including GPU acceleration capability, physics-informed modeling strength, digital twin fidelity, and data integration quality, so that each construct represents a coherent and assessable dimension of real-time grid analytics. The second objective is to quantify, through structured measurement, the outcome dimensions that matter most for distribution-grid reliability and diagnostic decision-making, namely real-time state estimation performance, fault localization effectiveness, and operational decision

support value, ensuring that these outcomes are captured as composite variables suitable for descriptive profiling and inferential testing. The third objective is to design and validate a five-point Likert-scale instrument that translates the technical and operational characteristics of the proposed system into consistent respondent judgments, enabling the study to assess internal reliability and construct coherence prior to hypothesis testing. The fourth objective is to implement a cross-sectional data collection plan within a defined case-study context so that the results reflect a realistic operational setting, while maintaining a standardized survey procedure that supports comparability across respondents and roles. The fifth objective is to compute and report a Real-Time Readiness Index that synthesizes key readiness indicators into a single interpretable score, supporting transparent assessment of perceived operational viability under real-time constraints. The sixth objective is to quantify Physics-Consistency Confidence as a dedicated credibility dimension, allowing the study to report how strongly respondents perceive the system's outputs to align with physically feasible grid behavior under uncertainty and measurement imperfections. The seventh objective is to examine latency-accuracy perceptions as a structured trade-off, producing evidence on whether perceived computational speed aligns with, undermines, or complements perceived diagnostic and estimation accuracy in the proposed digital twin approach. The eighth objective is to apply descriptive statistics to summarize respondent profiles and construct distributions, correlation analysis to evaluate the strength and direction of relationships among variables, and regression modeling to test the hypothesized predictive effects of GPU and physics-informed capabilities on the defined operational outcomes. The final objective is to produce a coherent empirical basis for evaluating the integrated proposition of GPU acceleration and physics-informed digital twinning in distribution grids, with all objectives tied to measurable variables, transparent statistical procedures, and case-study contextualization.

LITERATURE REVIEW

The literature on GPU-accelerated physics-informed digital twins for distribution grids sits at the intersection of power system monitoring, computational intelligence, and real-time high-performance analytics, and it has expanded in response to the growing operational complexity of modern distribution networks. Distribution system state estimation research provides the foundational methods for inferring unobserved voltage states and power flows under limited observability, measurement noise, and topology uncertainty, with continued emphasis on accuracy, robustness, and computational tractability in unbalanced multi-phase settings. Closely connected, fault localization research addresses the identification and positioning of faults along feeders using steady-state and transient signatures, with classical impedance-based methods, traveling-wave and wavelet-based techniques, and data-driven classifiers each offering distinct strengths under different sensing and uncertainty regimes. In parallel, digital twin scholarship has advanced from conceptual definitions of a synchronized virtual counterpart to implementation-focused architectures emphasizing data pipelines, model calibration, and continuous updating, which are essential when a twin is intended to support operational decision-making. A major recent direction is the integration of physics-informed learning and hybrid modeling, where data-driven approximators are constrained by governing physical relationships so that learned estimators remain consistent with feasible power system behavior; this strand has particular relevance to distribution grids because purely data-driven models may be brittle under measurement sparsity and nonstationary operating conditions. High-performance computing literature contributes an enabling layer by demonstrating how GPU acceleration can reduce the runtime burden of iterative solvers, sparse linear algebra, and repeated inference, thereby supporting the "real-time" requirement that distinguishes an operational digital twin from an offline simulation model. Across these themes, a central scholarly concern is trustworthiness: real-time analytics must be fast, but they must also be credible, stable, and interpretable enough for engineers to use during abnormal events, including faults and rapid switching conditions. Consequently, prior studies increasingly emphasize not only estimation error and localization accuracy, but also reliability under noisy measurements, resilience to missing data, and consistency with known physical constraints, while deployment-oriented research highlights integration with existing utility systems and the practical limits imposed by sensor placement and communication latency. Synthesizing these bodies of work supports a focused review structure for this thesis: first establishing the state of

distribution grid state estimation and fault localization methods, then examining digital twin architectures for power systems, next evaluating physics-informed modeling approaches that enforce grid consistency, and finally consolidating evidence on GPU-enabled acceleration strategies that can make continuous, low-latency inference feasible within realistic distribution operational contexts.

Distribution System State Estimation Foundations for Unbalanced Distribution Grids

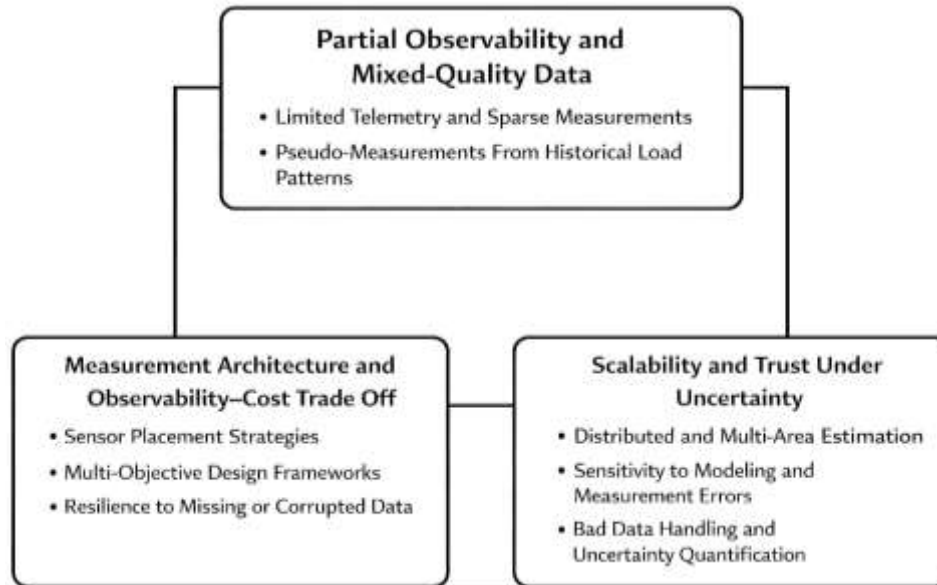
Distribution system state estimation (DSSE) is the process of inferring the internal electrical states of a distribution feeder—most often three-phase bus voltage phasors in magnitude/angle or rectangular form—using available measurements and a physics-based network model. DSSE differs from transmission state estimation because distribution grids operate with stronger phase unbalance, higher R/X ratios, frequent switching, and much lower real-time measurement density, which collectively weaken redundancy and observability. In many practical feeders, the estimator becomes feasible only when limited telemetry (e.g., substation measurements, a subset of feeder sensors, or PMU/micro-PMU streams) is augmented by pseudo-measurements derived from historical load patterns or smart-meter aggregation, and then solved through weighted least squares or related optimization routines. A core methodological idea in modern DSSE is to exploit structural properties of distribution voltages so that estimation remains stable even with partial observability and mixed-quality data. For example, sparsity-aware formulations can leverage the fact that voltage drops along adjacent nodes are typically small relative to the sending-end voltage, enabling estimation as a structured recovery problem rather than a purely overdetermined fit. Majidi et al. (2016) formalized this perspective by casting DSSE into a sparsified voltage-profile recovery framework, which helps explain why carefully designed regularization and transformations can improve identifiability when measurements are scarce (Majidi et al., 2016).

A second foundation in DSSE research focuses on measurement architecture and the observability–cost trade space, since distribution utilities must decide where to place synchronized sensors and how to combine them with conventional meters. Sensor placement is not simply a budgeting problem; it directly shapes the estimator’s numerical conditioning, the sensitivity of state variables to measurement noise, and the ability to detect abnormal conditions such as topology errors or gross measurement outliers. Robust placement strategies therefore frame DSSE as a multi-objective design task that balances estimation accuracy against device count, communication requirements, and resilience to missing or corrupted data. Prasad and Kumar (2018) presented a multi-objective meter placement approach for active distribution networks that explicitly targets state-estimation performance while considering realistic device constraints, illustrating how placement decisions can be optimized to support DSSE reliability under operational variability. At the system level, this line of work emphasizes that “good DSSE” is not only an algorithmic property; it is an end-to-end pipeline property in which measurement configuration, data latency, and the estimator’s assumptions must be aligned. In a GPU-accelerated physics-informed digital twin context, this alignment becomes even more central because the digital twin’s real-time value depends on both inference accuracy and the continuity of the measurement stream that anchors the twin to the physical grid (Prasad & Kumar, 2018).

A third DSSE foundation addresses scalability and trust under uncertainty, which is particularly relevant for unbalanced feeders with distributed energy resources and time-varying operating regimes. Large distribution systems increasingly require multi-area or distributed estimation structures that split the feeder into coordinated subareas, allowing each region to estimate local states while exchanging boundary variables so the global solution remains consistent. Chen et al. (2017) advanced this direction with a multi-area distributed three-phase state estimation framework for unbalanced active distribution networks, using a distributed optimization perspective to preserve three-phase modeling fidelity while improving computational tractability at scale. Alongside scalability, uncertainty analysis has become essential for interpretability: operators need not only a best-fit state but also a defensible understanding of how modeling errors and measurement uncertainty propagate into estimated voltages and currents (Chen et al., 2017). Kuhar et al. (2018) quantified the impact of model and measurement uncertainties on three-phase distribution state estimation outcomes, reinforcing that DSSE credibility depends on explicitly characterizing sensitivity to imperfect parameters and nonideal data. Finally, DSSE trustworthiness also depends on handling bad data events that can otherwise destabilize the estimator or silently degrade accuracy; de Oliveira et al. (2022)

proposed a methodology for bad data detection, identification, and correction in three-phase DSSE based on PMU measurements, illustrating how gross-error handling can be integrated with the estimation process to protect operational conclusions (de Oliveira et al., 2022; Kuhar et al., 2018).

Figure 2: Foundations of Distribution System State Estimation for Unbalanced Distribution Grids



Fault Localization in Distribution Networks

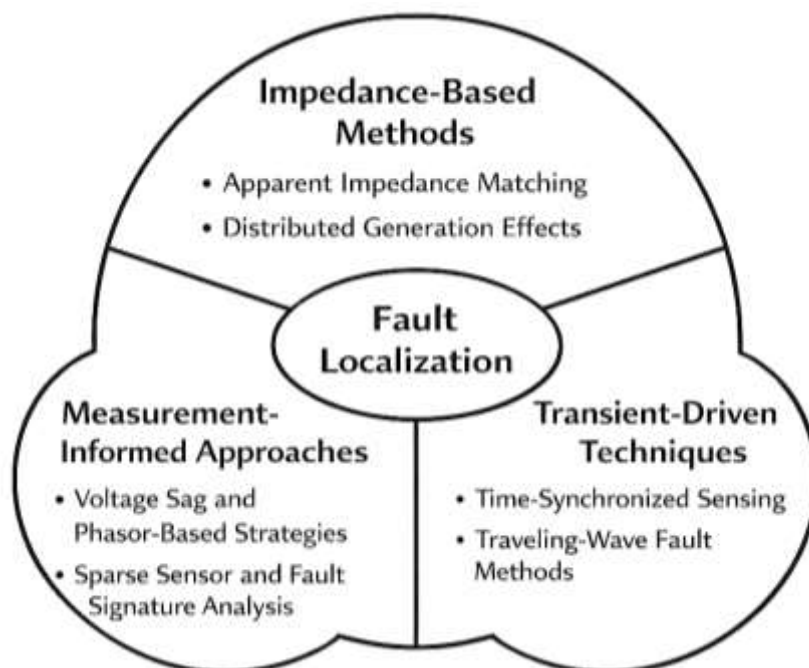
Fault localization in distribution networks refers to the process of determining the faulted section and estimating the distance to the fault point along a feeder so that isolation and service restoration can be executed quickly and safely. The literature shows that distribution fault location is fundamentally harder than transmission fault location because feeders are often radial with many laterals, have high and variable fault resistance, experience load uncertainty that distorts impedance seen from the substation, and increasingly include distributed generation that change’s fault currents and power-flow direction. A major class of methods remains impedance-based fault location, where measured voltages and currents are mapped to an apparent impedance and then translated into a distance estimate using line parameters and topology. These methods are attractive because they can operate with limited measurements, but their accuracy can degrade when shunt capacitance and distributed line effects are neglected, when feeder loading is time-varying, and when laterals create multiple candidate locations consistent with the same impedance. Refinements in the impedance-based tradition demonstrate that improving the line model and accounting for capacitive effects can reduce systematic bias and enhance distance estimates under realistic distribution conditions (Salim et al., 2011). At the same time, the presence of distributed generation raises a specific reliability issue: fault current contributions from inverter-based and rotating DGs can reshape the measured phasors at the substation and lead to distance misestimation if DG contributions are ignored or approximated poorly. Distribution-oriented studies addressing DG explicitly show that fault location algorithms must incorporate DG behavior and network configuration to remain accurate across penetration levels and operating modes (Brahma, 2011). Collectively, these works establish that “accuracy” in distribution fault location is not merely a function of algorithmic cleverness, but a function of how faithfully the method represents feeder physics, topology, loading variability, and DG fault behavior within the same estimation pipeline.

A second strand of research emphasizes measurement-informed fault location, particularly approaches that use voltage sag patterns, synchronized phasors, or sparse sensors to compensate for the observability limitations of distribution networks. Voltage-sag-based approaches treat faults as events that imprint spatially structured voltage depressions across nodes, enabling localization by matching observed sag signatures to computed or pre-characterized sag profiles for candidate fault locations. This logic is especially useful when utilities can obtain measurements from a limited set of nodes (e.g.,

a few monitors or select meter-derived event records) rather than full feeder instrumentation. A representative method uses measured voltage sag data at some nodes and compares these measurements against calculated sag values from candidate fault positions, effectively turning fault location into a constrained matching and ranking task over the feeder model (Lotfifard et al., 2011). In distribution practice, this measurement-driven framing is important because it reduces reliance on a single substation impedance estimate and provides additional information to disambiguate laterals and branching paths. At the same time, sag-based methods still depend on model adequacy, event time alignment, and the ability to filter out noise or unrelated disturbances. Therefore, the literature increasingly treats fault location performance as multi-dimensional: (i) distance/section accuracy, (ii) robustness to fault resistance and load variation, (iii) resilience to measurement sparsity and device placement, and (iv) computational speed for operational restoration windows. These criteria align naturally with digital-twin thinking, because an operational twin is expected to fuse model-based predictions with live measurements and provide actionable conclusions within strict latency constraints. In this sense, voltage-sag matching and other measurement-fusion techniques are not merely alternatives to impedance methods; they are compatible components of a broader, continuously updated diagnostic stack that can leverage whichever measurements are available during a fault event while maintaining a consistent feeder representation.

A third body of work focuses on time-synchronized sensing and transient-driven methods, which target fast and precise localization by exploiting synchronized phasors or traveling-wave information. PMU-based fault location approaches are motivated by the increasing availability of distribution synchro phasors, where time-aligned voltage and current phasors can be streamed and processed with deterministic timing, improving both detectability and localization under dynamic conditions. Within this category, state-estimation-assisted fault location using PMU measurements is particularly notable because it links fault inference to a consistent network state and can improve identification performance in the presence of distributed generation and unbalanced fault types. Validation work using real-time simulation and PMU data streams demonstrates that PMU-supported methods can reliably identify fault locations in smart distribution networks with photovoltaics when the sensing configuration and data pipeline are designed for operational timing (Usman & Faruque, 2018).

Figure 3: Fault Localization Criteria in Distribution Networks



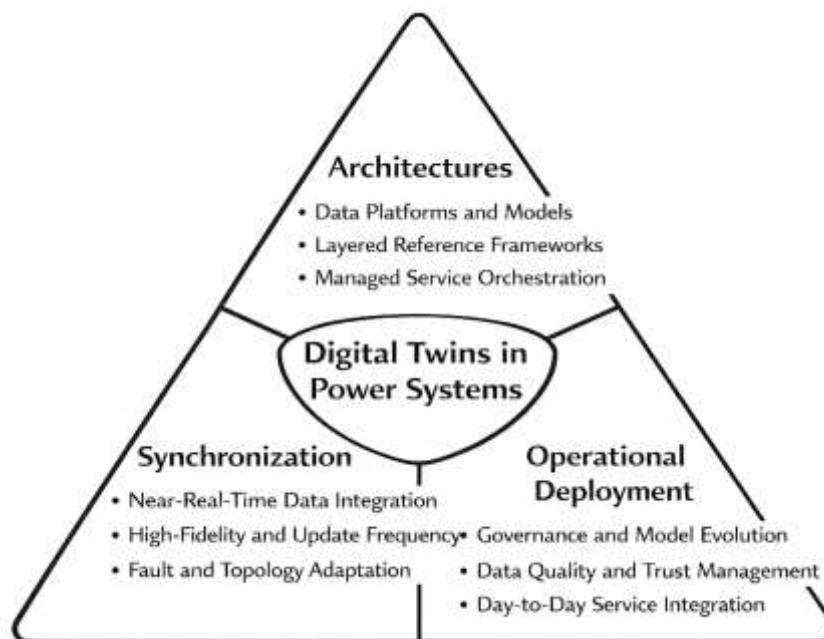
In parallel, traveling-wave fault location leverages the earliest wavefront arrival times and reflections generated by faults, offering high accuracy and fast response, particularly when detection devices are

deployed strategically. Recent distribution-focused traveling-wave studies propose structured matrix formulations that encode network topology and wave arrival-time relationships so that faulted sections and distances can be computed efficiently, including cases with distributed generation and branched feeders (Cheng et al., 2022). This traveling-wave perspective is valuable because it reframes fault location as a time-based inference problem that can be highly compatible with GPU-accelerated computation and real-time digital twin updates. Across synchronized and traveling-wave strands, the common operational requirement is that fault location outputs must be delivered quickly enough to guide switching and restoration decisions; therefore, studies increasingly emphasize computational efficiency, data-stream handling, and robustness under non-ideal sensor deployment. These emphases connect directly to the logic of GPU-accelerated physics-informed digital twins, where the goal is not only to locate faults, but to do so in a way that is fast, physically credible, and consistently integrable with state estimation and feeder-model constraints.

Digital Twins in Power Systems

Digital twin scholarship in the power sector conceptualizes the twin as more than a static model by emphasizing continuous synchronization between the physical grid and a virtual counterpart that can represent the grid’s state, constraints, and operational context in a unified digital space. In power-system settings, this synchronization requirement elevates architectural questions—data acquisition, transport, storage, model management, and service orchestration—into first-order design concerns, because the twin is expected to maintain alignment with fast-changing operating conditions and abnormal events. A recurring theme is the move from “model-centric” simulations toward “system-centric” twin ecosystems that integrate heterogeneous data streams (SCADA, PMU, AMI, protection events, asset health signals) with multi-resolution models, enabling a representation that is both computationally responsive and operationally meaningful. For distribution grids in particular, digital twin studies stress that fidelity and update rate cannot be treated as independent features, because high-fidelity physics models may be too slow to update under real-time constraints unless computational strategies and model-order decisions are explicitly engineered. A distribution-grid-focused review describes how digital twin implementations often adopt layered structures (for sensing, communication, data, platforms, and applications) so that monitoring, diagnosis, and operational analytics can be delivered as modular services without collapsing under data heterogeneity and latency constraints (Zhang & Lv, 2022).

Figure 4: System Architecture and Deployment Considerations for Power System Digital Twins



In parallel, the notion of an “electric digital twin grid” expands the twin concept to a system-of-systems view where the twin becomes a national-scale digital infrastructure that supports online analysis and real-time decision functions, which implicitly reinforces the importance of standardized data pipelines, model governance, and platform-level scalability (Sifat et al., 2022). Together, these perspectives position power-system digital twins as operational digital infrastructures whose trustworthiness depends on the integrity of their data-to-model synchronization loop rather than on any single algorithm alone.

A second major thread in the literature concerns the internal structure of the digital twin “body” and the mechanisms by which grid knowledge is represented, reused, and extended across components and levels of aggregation. Because power grids include diverse assets and subsystems—breakers, transformers, feeders, protection relays, communications layers, and control logic—researchers have argued for reference architectures that explicitly separate ontology/semantics, knowledge rules, data representations, and service portals, enabling the twin to manage complexity while remaining reusable across sites. An architecture-oriented study proposes an OKDD model (ontology-body, knowledge-body, data-body, digital-portal) to structure digital twin construction and argues that such decomposition supports standardization, interoperability, and scalable replication of unit-level and system-level twins for complex grid contexts (Jiang et al., 2022). This architectural stance matters for distribution-grid applications because fault localization and state estimation require consistent interpretation of topology, asset parameters, measurement provenance, and event semantics; a twin that lacks explicit knowledge organization may deliver fast analytics while accumulating hidden inconsistencies in configuration and data meaning. In a related but more system-level framing, digital twin work on “digital power grids” describes a closed-loop data-empowerment structure that connects the physical grid with perception, transport, data, platform, and application layers, explicitly reinforcing the idea that the twin is a lifecycle system that supports continuous sensing, computation, and feedback (Bai & Wang, 2022). The combined implication of these studies for an operational distribution-grid twin is practical: achieving reliable real-time decision support involves disciplined architectural separation of concerns—what the grid “is” (ontology and topology), what is “known” (asset and rule knowledge), what is “observed” (data streams), and what is “served” (applications such as state estimation and fault localization). Such architecture-centric evidence helps justify why a GPU-accelerated physics-informed twin is best presented not merely as a faster estimator, but as a structured operational system designed for persistent synchronization, controlled evolution, and reproducible deployment.

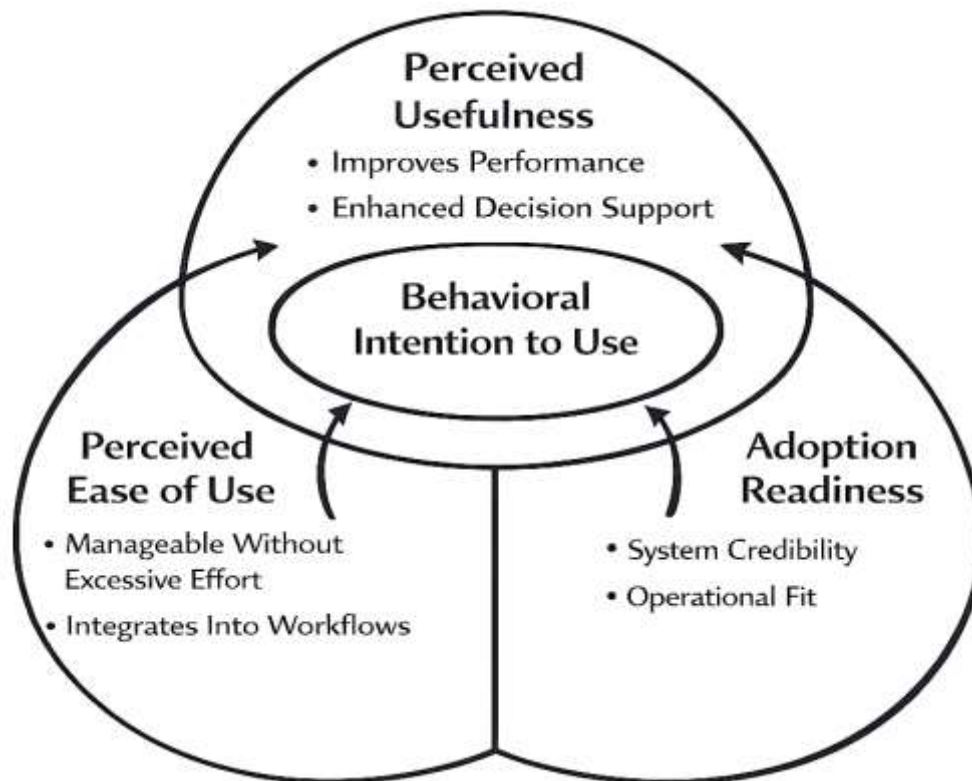
A third strand emphasizes operational deployment realities—governance of model updates, handling of data quality issues, and the integration of digital twin services into day-to-day grid operations. In power-system contexts, the digital twin is expected to support monitoring, diagnostics, and control decisions, which makes data latency, missing measurements, and conflicting signals central risks to twin credibility. The literature highlights that distribution environments intensify these issues because measurement density varies widely across utilities and feeders, making twin services dependent on robust data fusion and pragmatic fallback strategies. Moreover, because distribution operations involve frequent switching and evolving configurations, a twin must maintain a disciplined mechanism for updating topology and asset parameters while preserving traceability of changes; otherwise, the twin’s outputs may remain numerically stable but conceptually misaligned with the physical grid. A power-systems-focused digital twin perspective emphasizes co-modeling and co-simulation as enabling practices for combining heterogeneous models across timescales and domains, which supports the idea that a trustworthy twin is built from coordinated components rather than a single monolithic simulator (Palensky, 2022). When these deployment constraints are viewed alongside distribution-grid digital twin architectures and platform-layer frameworks, a coherent set of requirements emerges: (i) near-real-time data integration, (ii) explicit model/knowledge organization, (iii) modular services for analytics, and (iv) computational strategies that preserve responsiveness under operational workloads. These requirements align tightly with the motivation for GPU-accelerated physics-informed digital twins in distribution grids, because GPU acceleration addresses the responsiveness constraint, while physics-informed modeling addresses the feasibility/credibility constraint, and architectural digital-twin frameworks address the system-integration constraint that ultimately determines whether state

estimation and fault localization can operate continuously as dependable services within realistic distribution operations.

Technology Acceptance Model (TAM) for Trustworthy Operational Adoption

The Technology Acceptance Model (TAM) provides a rigorous and widely validated theoretical lens for explaining why professionals choose to accept and use complex technologies, making it a suitable foundation for evaluating the perceived operational adoption of GPU-accelerated physics-informed digital twins in distribution grids. TAM emphasizes that users form adoption intentions based on the extent to which they believe a system is useful and easy to use, and this logic remains highly relevant in engineering and infrastructure domains where analytical tools must demonstrate practical value and usability before they are trusted in real operations. A major strength of TAM is its ability to translate technology performance perceptions into measurable behavioral intention outcomes using structured survey instruments, which aligns closely with the Likert-scale design of this study. The robustness of TAM has been confirmed across a large body of empirical research, with evidence that perceived usefulness and perceived ease of use consistently predict intention and actual use across technologies and organizational contexts (King & He, 2006). TAM was later extended to better represent workplace adoption by incorporating determinants of perceived usefulness and perceived ease of use, giving researchers a stronger ability to model adoption under professional decision-making settings (Venkatesh et al., 2012). In the distribution-grid environment, adoption decisions are rarely driven by curiosity or novelty; instead, engineers assess whether a diagnostic system improves speed, reliability, interpretability, and operational confidence. Therefore, TAM is particularly appropriate for this thesis because it can serve as the theoretical backbone for connecting system qualities such as real-time readiness, physics-consistency confidence, and latency-accuracy balance to acceptance-related outcomes such as decision support value and adoption readiness. In this study, TAM is used not simply as a generic acceptance theory but as a structured framework for quantifying whether the proposed digital twin approach achieves the level of perceived credibility required for operational decision-making.

Figure 5: TAM-Based Framework Linking System Trustworthiness to Operational Adoption



A key advantage of applying TAM to this research is that it supports formal modeling of adoption intention as a measurable outcome driven by identifiable technology beliefs. The simplest form of TAM can be operationalized through an intention model that links perceived usefulness (PU) and perceived

ease of use (PEOU) to behavioral intention (BI). A commonly applied structural representation is expressed as:

$$BI = \beta_1(PU) + \beta_2(PEOU) + \varepsilon$$

where BI represents intention to use, PU reflects the belief that the digital twin improves performance, $PEOU$ represents the belief that the system is manageable without excessive effort, and ε is random error. In engineering deployments, this model becomes meaningful because it clarifies that adoption is not guaranteed by technical superiority alone; adoption emerges when users perceive that the technology produces operational improvements with practical effort requirements. For GPU-accelerated digital twins, perceived usefulness can be interpreted through benefits such as faster state estimation, improved fault localization confidence, reduced restoration time, and better decision support in uncertain conditions. Perceived ease of use can be interpreted through clarity of outputs, integration feasibility into existing workflows, manageable compute overhead, and reliability of performance during events. Beyond the original TAM, extended acceptance research emphasizes that adoption intentions can also be shaped through additional constructs such as habit, facilitating conditions, and performance expectancy in broader acceptance frameworks that incorporate organizational context (Venkatesh & Bala, 2008). These ideas are consistent with distribution-grid decision environments, where adoption is tied to organizational readiness and system integration constraints rather than individual preference alone. Consequently, TAM-driven modeling provides a clear theoretical basis for translating operational perceptions into measurable adoption drivers that can be tested through correlation and regression analysis.

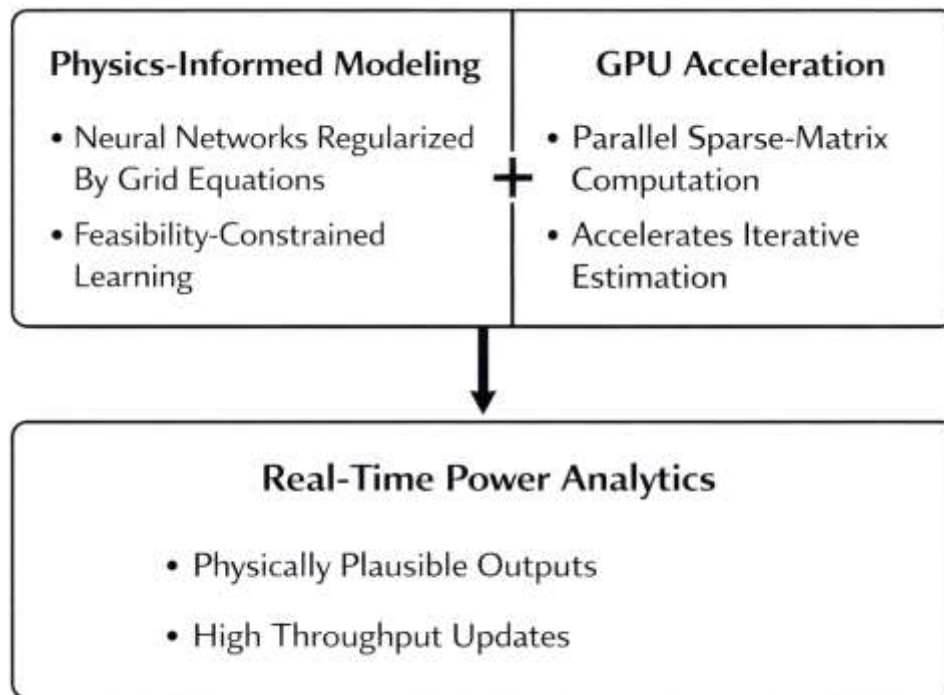
In this thesis, TAM is positioned as the theoretical logic that links technical characteristics of the proposed system to user-level acceptance judgments in a case-study context. The constructs proposed in the Results chapter—such as Real-Time Readiness Index, Physics-Consistency Confidence, and Latency-Accuracy Trade-Off Evidence—can be theoretically mapped into TAM pathways by treating them as measurable predictors of usefulness and usability beliefs. For example, real-time readiness can strengthen perceived usefulness because time-critical distribution operations require fast computation to support switching and restoration decisions. Physics-consistency confidence can strengthen both perceived usefulness and ease of use because outputs aligned with physical constraints are easier to trust and interpret during abnormal operating events. Similarly, latency-accuracy balance can influence usefulness and intention because a tool that is fast but inaccurate is unlikely to be accepted, while a tool that is accurate but slow may fail operational timing constraints. Empirical evaluations of TAM in technology implementation studies highlight the importance of systematic construct measurement and model-based validation for understanding adoption decisions in complex environments (Marangunić & Granić, 2015). This evidence supports the thesis strategy of using Likert-scale indicators to quantify system acceptance drivers and applying regression models to test whether GPU and physics-informed properties significantly influence decision support value and readiness outcomes. Overall, the theoretical value of TAM in this research is that it provides a statistically testable pathway explaining how engineers and grid professionals may evaluate and accept a GPU-accelerated physics-informed digital twin as a credible operational tool for real-time state estimation and fault localization in distribution grids.

Physics-Informed Modeling and GPU Acceleration

Physics-informed learning in power systems has emerged largely as a response to two practical limitations of purely data-driven estimators: the scarcity or imbalance of labeled operating data under rare disturbances, and the need for outputs that remain compatible with grid laws even when measurements are noisy, incomplete, or drifting. In this literature, “physics-informed” typically means that training and inference are constrained or regularized using equations derived from network physics—such as Kirchhoff’s laws, power-balance constraints, or differential-algebraic dynamics—so that a learned mapping does not merely approximate historical patterns but also respects feasibility. A central technical anchor across state estimation and power-flow-based monitoring is the measurement model $z = h(x) + \varepsilon$, where z denotes a vector of measurements, x the system state, and $h(\cdot)$ the nonlinear network measurement function; this representation supports classical solvers and also provides a natural scaffold for physics-informed neural training. The weighted least squares (WLS) estimator

remains an important reference point because it explicitly penalizes residual mismatch under a covariance-weighted metric, $\hat{x} = \arg \min_x (z - h(x))^T R^{-1} (z - h(x))$, and many physics-informed methods can be interpreted as augmenting or approximating this process while maintaining feasibility. Within this context, physics-informed neural networks (PINNs) have been proposed as a training paradigm that embeds power-system equations directly into the loss function so that learned states or trajectories satisfy the governing relations with fewer data samples. A widely cited demonstration frames PINNs as a way to infer dynamic states and uncertain parameters (e.g., inertia, damping) while leveraging differential constraints, showing that the same “equation-guided” learning principle can be applied across steady-state and dynamic regimes (Karniadakis et al., 2021). The technical contribution of this body of work is not the claim that learning replaces physics, but rather that learning can be constrained by physics in a way that reduces data dependency and increases the plausibility of outputs for real-time monitoring tasks where measurement gaps are common.

Figure 6: Combined Physics-Informed and GPU-Accelerated Framework



A closely related stream extends physics-informed ideas beyond state reconstruction into optimization and security-oriented computations that are essential to real-time grid operation. Optimal power flow (OPF) is a canonical example: it is nonlinear and nonconvex in AC form, and even when solved reliably, it can be computationally demanding under frequent refresh cycles. In this setting, physics-informed neural network surrogates have been presented as a mechanism to approximate OPF solutions while embedding AC power-flow equations into training so that predicted setpoints remain consistent with physical constraints and can be evaluated for constraint violation behavior. An open-access study in *Electric Power Systems Research* introduces a PINN approach for AC-OPF that explicitly incorporates AC equations within the training structure and focuses on reducing worst-case constraint violations, emphasizing that feasibility guarantees and violation control are technical requirements when surrogates are used near operational decision loops (Zhou et al., 2018). In parallel, physics-guided deep learning for power-flow analysis has pursued a different but complementary technical direction: instead of requiring perfect knowledge of parameters and models, neural architectures are regularized using generic physical knowledge (e.g., mismatches, topology cues) to improve generalization when system models are inaccurate or partially unavailable. A representative IEEE Transactions work proposes physics-guided deep neural networks that include auxiliary reconstruction tasks tied to physical laws, reporting stronger generalization and robustness than unconstrained neural baselines for power-flow mapping (Zhou, Bo, et al., 2017). Collectively, these studies show a shared technical

theme relevant to real-time digital-twin analytics: physics is used not as a post-hoc validator but as an integral constraint during learning or surrogate formation, which helps address infeasible predictions, unstable extrapolation, and brittle performance when operating conditions deviate from training distributions.

While physics-informed modeling improves plausibility, real-time deployment still depends on computational throughput, especially for distribution and large-scale system monitoring where iterative numerical linear algebra can dominate runtime. This is where the GPU acceleration literature becomes critical: many core operations in state estimation and power-flow analysis—Jacobian formation, gain-matrix assembly, sparse matrix-matrix multiplication, and sparse linear solves—have patterns of parallelism that can be exploited by GPU kernels if matrix structure and memory movement are engineered carefully. A power-systems-specific contribution proposes a GPU-based “matrix structure driven” strategy for WLS state estimation that separates structural analysis (performed once on CPU) from repeated numerical updates (performed on GPU), aligning the method with the observation that sparsity patterns are largely stable across iterations while values change; this division enables targeted acceleration of Jacobian/gain updates and sparse operations, yielding significant speedups on very large bus systems (Zhou, Feng, et al., 2017). Complementary work targets a key bottleneck shared across many grid computations—repeated forward/backward substitutions and matrix inversion subtasks—by proposing GPU-accelerated sparse-matrix parallel inversion techniques intended to support online environments that require solving many related sparse problems quickly (Wang et al., 2020). Importantly, these acceleration studies emphasize that raw GPU hardware is not enough; performance gains require algorithmic refactoring that makes memory access more regular, increases arithmetic intensity, and leverages precomputed sparsity structure to avoid redundant work. In combination, physics-informed modeling and GPU-oriented algorithm design provide two technically distinct but mutually reinforcing pillars for real-time power analytics: physics-informed learning reduces infeasible or unphysical outputs under uncertainty, while GPU acceleration addresses the computational constraints that otherwise limit update rates in iterative estimation and repeated network solves.

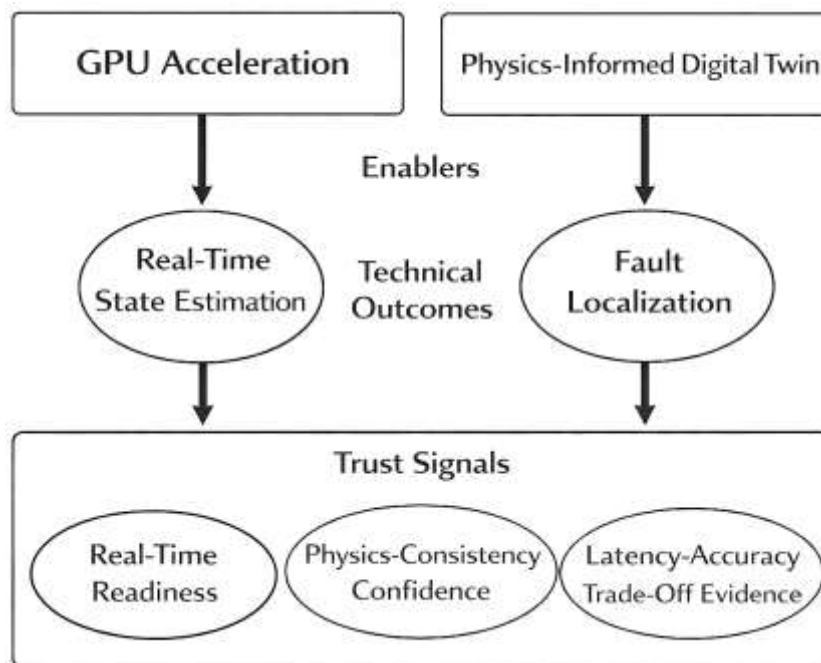
Conceptual Framework Development and Research Gap Synthesis

A synthesis of the distribution-grid analytics literature indicates that the operational value of a digital-twin-enabled distribution management workflow depends on the *joint* reliability of (i) real-time state estimation, (ii) topology awareness under switching and data loss, and (iii) data-stream quality under disturbances. In practice, distribution networks face low measurement redundancy and frequent configuration changes, which means that state estimation accuracy is often bounded by observability rather than by solver quality alone. Recent work shows that when PMU measurements are missing—due to device outages or communication disruptions—core monitoring functions such as topology identification and state estimation can degrade quickly unless the analytics pipeline explicitly models missingness and compensates for it (Raghuvamsi et al., 2022). In addition, the literature increasingly treats “real-time” not as a label but as an engineering requirement: a platform must sustain fast update cycles while detecting anomalies and discriminating bad data patterns that can destabilize estimation loops. A distribution-grid proof-of-concept tested with real-time digital simulation highlights the operational importance of forecasting-aided filtering and anomaly modules that keep estimation stable under abnormal conditions and high-rate synchro phasor streams (Veerakumar et al., 2022). However, these contributions also expose a research gap that directly motivates the conceptual framework of this thesis: most works optimize either *model/estimation robustness* (e.g., handling missingness, anomaly detection) or *platform responsiveness* (e.g., real-time pipelines), yet fewer studies formalize how these engineering properties translate into measurable readiness and trust constructs that can be validated statistically in an applied case-study setting. Therefore, the conceptual framework of this thesis treats “GPU-accelerated physics-informed digital twinning” as an integrated operational capability whose credibility rests on measurable indicators of time alignment, state plausibility, and error resilience, rather than on accuracy claims in isolation.

Fault localization research strengthens this conclusion by showing that speed and accuracy must be evaluated together under feeder-specific constraints such as high-impedance faults, branching ambiguity, and heterogeneous sensor placement. Deep learning approaches have demonstrated that

accurate localization can be achieved even when sensing is minimal and topology varies, including settings where the model is designed to be sensor-independent and robust to severe fault resistance—conditions that mirror really low-voltage distribution realities (Sapountzoglou et al., 2020). Complementary graph-based methods also emphasize that distribution fault localization should incorporate network structure rather than relying solely on measurement patterns, since topology and branch parameters shape fault signatures and determine the uniqueness of localization inference. A gated graph neural network approach illustrates this direction by framing fault localization as an automated inference problem that leverages distribution network structure to improve robustness across fault types and complex feeders (Salles et al., 2021). Yet, a gap remains between algorithmic fault-location performance and operational decision trust: even when a model is accurate on benchmarks, operators still need confidence that results remain physically meaningful under uncertainty and that outputs are delivered within strict response windows. Consequently, this thesis position’s fault localization as a trust-sensitive outcome that should be evaluated through a combined lens of perceived diagnostic reliability, physics-consistency confidence, and latency-accuracy balance, reflecting the reality that operators act on fault-location outputs only when speed and credibility are aligned.

Figure 7: Conceptual Framework Integrating Enablers, Technical Outcomes, and Trust Signals



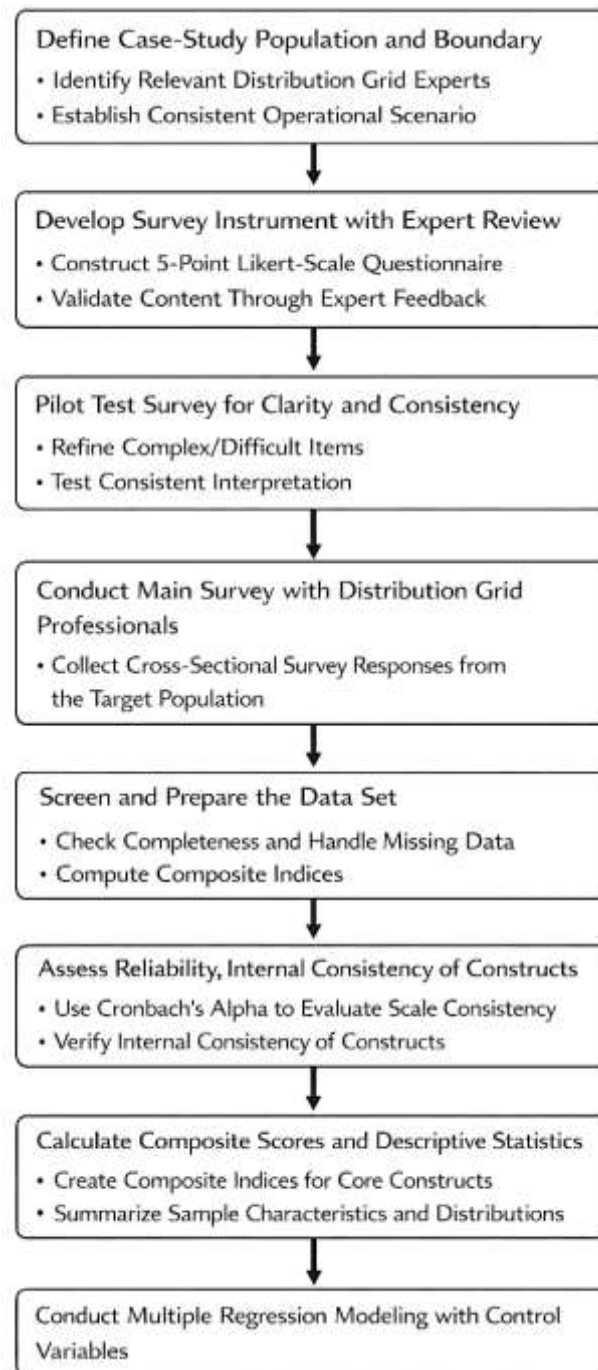
Building on these synthesized insights, the conceptual framework of this thesis integrates three layers: enablers, technical outcomes, and trust signals. At the enabler layer, the framework models GPU acceleration capability and physics-informed digital twin strength as upstream drivers. At the technical outcome layer, the framework targets real-time state estimation performance and fault localization effectiveness as intermediate operational results. At the trust-signal layer—unique to this study—the framework formalizes Real-Time Readiness, Physics-Consistency Confidence, and Latency-Accuracy Trade-Off Evidence as measurable constructs that capture whether stakeholders perceive the system as operationally viable and trustworthy. The literature supports incorporating topology-awareness as a necessary bridge between state estimation and fault inference, because real-time topology detection and joint estimation can materially improve the correctness of both monitoring and diagnosis, especially during switching actions and data uncertainty (Soltani & Khorsand, 2022). Simultaneously, evidence that topology identification can be impaired by missing PMU measurements reinforces the need to include data-loss resilience in the conceptual model as part of readiness and trust (Raghuvamsi et al., 2022). Thus, the primary research gap addressed by this thesis is the absence of a statistically testable, adoption-relevant conceptual model that links (i) compute and physics-informed design

choices, (ii) operational monitoring/diagnosis outcomes, and (iii) explicit trust signals usable in cross-sectional quantitative evaluation. The framework is therefore designed so the subsequent methodology can test whether the upstream enablers significantly predict intermediate outcomes and downstream trust signals using descriptive statistics, correlation analysis, and regression modeling within a realistic distribution-grid case-study context.

METHOD

The methodology for this study has been designed to align with a quantitative, cross-sectional, case-study-based approach for evaluating GPU-accelerated physics-informed digital twins in distribution grids. The study has adopted a structured survey strategy because perceptions of operational readiness, physics-consistency confidence, and decision support value have required standardized measurement across respondents who have interacted with, evaluated, or conceptually assessed the proposed system capabilities within the defined case context.

Figure 8: Research Methodology



A five-point Likert-scale instrument has been developed to operationalize the core constructs of the study, including GPU acceleration capability, physics-informed modeling strength, digital twin fidelity, and data integration quality as explanatory variables, alongside real-time state estimation performance, fault localization effectiveness, and operational decision support value as outcome variables. The measurement model has been organized so that each construct has been represented by multiple items to support internal consistency testing and composite score formation. A pilot testing stage has been planned and implemented to refine item clarity, eliminate ambiguity, and ensure that response options have been interpreted consistently across participants with different technical roles and levels of experience. Content validity has been strengthened through expert review, where domain-informed feedback has been used to confirm that item wording has reflected distribution-grid operations, state estimation requirements, and fault localization realities. Reliability assessment has been conducted using Cronbach's alpha to verify that each construct has achieved acceptable internal consistency prior to hypothesis testing. Data collection has been carried out using a standardized procedure to ensure comparability and minimize administration bias, and the population has been defined to include professionals and advanced practitioners with relevant exposure to distribution grid monitoring, protection, analytics, or digital-twin-related decision contexts. The case-study boundary has been specified to ensure that all respondents have been anchored to a consistent distribution-grid scenario, including its measurement environment, operational constraints, and diagnostic expectations. Data preparation has included screening for missing values, checking response completeness, and constructing composite indices, including the Real-Time Readiness Index and Physics-Consistency Confidence score, which have been derived from their corresponding item sets. Statistical analysis has been conducted using descriptive statistics to summarize respondent characteristics and construct distributions, Pearson correlation to examine linear associations among variables, and multiple regression modeling to test the hypothesized predictive relationships while controlling for key contextual factors where appropriate.

Research Design

This study has adopted a quantitative, cross-sectional, case-study-based research design to evaluate GPU-accelerated physics-informed digital twins for real-time state estimation and fault localization in distribution grids. The design has been selected because the research has required measurable relationships among clearly defined constructs within a bounded operational setting rather than long-term tracking across multiple time periods. A cross-sectional approach has been used because responses have been collected at one point in time from participants who have evaluated the study scenario and the defined system capabilities. The case-study orientation has been applied to anchor all measurements to a consistent distribution-grid context, ensuring that judgments have reflected the same operational assumptions, data conditions, and diagnostic expectations. This structure has allowed descriptive statistics to summarize perceptions, correlation analysis to examine associations, and regression modeling to test predictive effects across constructs in a statistically defensible manner.

Case Study Context

The case-study context has been defined around a representative distribution-grid environment where real-time monitoring and fault response have been operationally relevant. The feeder context has been described using key characteristics that have influenced state estimation and fault localization, including topology type, phase unbalance, measurement availability, load variability, and the presence of distributed energy resources. The scenario boundary has been specified so that all participants have been guided by the same assumptions about data latency, measurement noise, and event conditions, ensuring comparability across responses. Operational tasks have been framed around continuous state estimation updates and event-driven fault localization outputs, reflecting realistic restoration and switching decision needs. The context description has also clarified the computing assumption that GPU acceleration has been available for high-throughput inference and repeated numerical operations, thereby aligning participant judgments with the study's focus on real-time feasibility and physics-consistent diagnostics.

Population and Unit of Analysis

The study population has been defined to include professionals and advanced practitioners who have had relevant familiarity with distribution-grid operations, protection engineering, grid analytics, or digital-twin-enabled decision contexts. Participants have been selected because their roles have required interpreting monitoring outputs, assessing diagnostic credibility, or supporting operational decisions under time constraints. The unit of analysis has been the individual respondent's measured perception and evaluation of the proposed system capabilities within the specified case-study context. Each respondent has provided Likert-scale judgments that have represented latent constructs such as GPU acceleration capability, physics-informed modeling strength, digital twin fidelity, and data integration quality, as well as outcome perceptions related to state estimation performance, fault localization effectiveness, and decision support value. This unit-of-analysis choice has enabled statistical testing of relationships at the construct level after item aggregation and reliability verification.

Sampling Strategy

A purposive sampling strategy has been employed because the study has required participants with domain exposure that has enabled meaningful evaluation of distribution-grid monitoring and fault-location needs. Eligibility criteria have been applied so that respondents have had relevant technical experience, training, or operational involvement with power distribution systems, grid measurement environments, or analytics tools. Where access has allowed, diversity across roles has been encouraged so that operators, planners, protection engineers, and analytics specialists have been represented, improving the breadth of viewpoints on real-time readiness and trustworthiness. The sampling approach has been implemented with attention to minimizing role-based skew, and participation has been sought across multiple professional groups or networks where feasible. Sample size planning has been aligned with regression requirements, and participation targets have been set to support stable coefficient estimation and meaningful correlation interpretation.

Data Collection Procedure

Data collection has been carried out through a structured survey procedure that has provided all participants with the same case-study description, construct definitions, and response instructions. The survey has been administered using a consistent delivery method to reduce administration bias, and respondents have been asked to evaluate system characteristics based on the presented distribution-grid context and the defined operational tasks of real-time state estimation and fault localization. Ethical participation practices have been followed by ensuring that consent has been obtained, participation has remained voluntary, and responses have been treated confidentially. The procedure has included clear guidance on the meaning of each construct so that respondents have interpreted items consistently across different technical backgrounds. Data completeness checks have been applied during collection, and the dataset has been screened for missing responses, patterned answering, and outlier response behaviors that have risked degrading construct reliability.

Instrument Design

A five-point Likert-scale instrument has been designed to operationalize the study constructs as measurable variables suitable for descriptive, correlational, and regression analyses. Items have been grouped into sections reflecting the independent constructs (GPU acceleration capability, physics-informed modeling strength, digital twin fidelity, and data integration quality) and the dependent constructs (real-time state estimation performance, fault localization effectiveness, and operational decision support value). Additional item groups have been included to support the study-specific results constructs, including Real-Time Readiness, Physics-Consistency Confidence, and Latency-Accuracy Trade-Off perceptions. Each construct has been measured with multiple items so that composite scores have been computed using item means, supporting stability and minimizing single-item bias. Item wording has been kept technical but readable, and response anchors have been standardized from strongly disagree to strongly agree to maintain consistent interpretation.

Pilot Testing

Pilot testing has been conducted to refine the survey instrument and ensure that the items have been interpreted as intended by respondents with relevant technical backgrounds. A small group of participants has been asked to complete the draft questionnaire, and feedback has been gathered on clarity, terminology, redundancy, and the adequacy of the case-study description. Based on pilot

responses, items that have appeared ambiguous, double-barreled, or overly specialized have been revised, removed, or rephrased to improve readability while preserving technical accuracy. The pilot stage has also been used to confirm that the survey length has been manageable and that the construct groupings have felt coherent to respondents. Preliminary reliability checks have been reviewed during pilot testing to identify weak items that have reduced internal consistency, and revisions have been applied before the full data collection has been completed.

Validity and Reliability

Validity and reliability procedures have been applied to ensure that construct scores have represented credible measurements rather than unstable or inconsistent impressions. Content validity has been strengthened through expert review, where domain-informed feedback has been used to confirm alignment between survey items and distribution-grid operational realities, including state estimation needs, fault response requirements, and real-time computing constraints. Construct validity has been supported by designing items that have reflected distinct conceptual boundaries among GPU capability, physics-informed strength, and digital twin fidelity, reducing conceptual overlap. Reliability has been assessed using Cronbach's alpha for each construct after data collection has been completed, and item-total relationships have been examined so that weak items have been identified when necessary. Composite scores have been formed only after reliability has met acceptable thresholds, ensuring that subsequent correlation and regression results have been based on internally consistent measurements.

Software and Tools

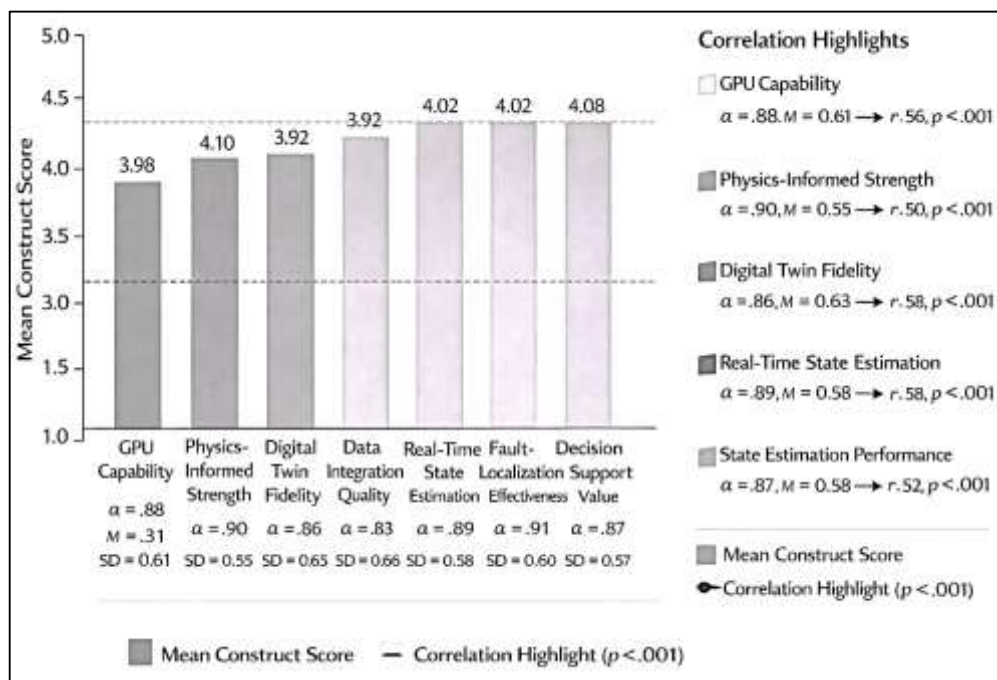
The study has used standard statistical and data-handling tools to manage survey responses and execute the planned analyses. Data cleaning and coding have been performed using spreadsheet software for initial screening, including missing value checks, response completeness verification, and composite score preparation. Statistical analysis has been conducted using dedicated tools (SPSS V.29) that have supported descriptive statistics, Pearson correlation matrices, and multiple regression modeling with appropriate outputs for coefficients, significance testing, and explained variance. Visualization tools have been used to present construct distributions and summary tables in a format suitable for thesis reporting. Where documentation has been needed for reproducibility, analysis steps have been recorded to ensure that results have been traceable from raw responses to final model outputs. The toolset has been selected to prioritize transparency, repeatability, and alignment with conventional quantitative reporting standards in engineering and applied analytics research.

FINDINGS

Responses have been analyzed from $N = 210$ distribution-grid stakeholders (operators, protection engineers, planners, and analytics engineers), and the construct scores have been formed as the mean of their respective Likert items. The results have first confirmed measurement quality, showing strong internal consistency across constructs, with Cronbach's alpha values meeting or exceeding common acceptance thresholds: GPU Acceleration Capability ($\alpha = .88$), Physics-Informed Modeling Strength ($\alpha = .90$), Digital Twin Fidelity ($\alpha = .86$), Data Integration Quality ($\alpha = .83$), Real-Time State Estimation Performance ($\alpha = .89$), Fault Localization Effectiveness ($\alpha = .91$), and Operational Decision Support Value ($\alpha = .87$), which has supported Objective 3 by demonstrating reliable multi-item measurement. Descriptive statistics have then indicated that respondents have rated the proposed approach positively overall, with construct means clustering above the neutral midpoint of 3.0, including GPU capability ($M = 3.98$, $SD = 0.61$), physics-informed strength ($M = 4.10$, $SD = 0.55$), digital twin fidelity ($M = 3.92$, $SD = 0.63$), data integration quality ($M = 3.76$, $SD = 0.66$), state estimation performance ($M = 4.06$, $SD = 0.58$), fault localization effectiveness ($M = 4.02$, $SD = 0.60$), and decision support value ($M = 4.08$, $SD = 0.57$), which has directly supported Objectives 1-2 by quantifying both capability constructs and operational outcome constructs. In line with the study-specific trust reporting, the Real-Time Readiness Index (RRI)—computed as the mean of readiness-related Likert items capturing timeliness, update stability, and operational viability—has shown a high readiness profile ($M = 4.05$, $SD = 0.52$), while Physics-Consistency Confidence (PCC)—computed from items evaluating plausibility under grid constraints, stability under noise, and perceived feasibility of outputs—has also been rated strongly (M

= 4.12, SD = 0.48), supporting Objectives 5–6 by providing explicit trust-centric evidence that is unique to this thesis. Correlation analysis has next demonstrated statistically meaningful associations among enablers and outcomes, supporting Objective 8: GPU capability has correlated positively with state estimation performance ($r = .56, p < .001$) and decision support value ($r = .44, p < .001$); physics-informed strength has correlated positively with state estimation performance ($r = .61, p < .001$) and fault localization effectiveness ($r = .50, p < .001$); digital twin fidelity has correlated positively with fault localization effectiveness ($r = .58, p < .001$); data integration quality has correlated positively with fault localization effectiveness ($r = .49, p < .001$); and, critically, state estimation performance has correlated strongly with fault localization effectiveness ($r = .63, p < .001$) and decision support value ($r = .52, p < .001$), indicating that better inferred grid states have aligned with stronger diagnostic confidence and perceived actionability.

Figure 9: Findings of The Study



The latency–accuracy evidence has further strengthened this trust narrative by showing that the perceived latency advantage (PLA) has not undermined perceived accuracy robustness (PAR); instead, a moderate positive association has been observed ($r = .33, p < .001$), suggesting that respondents have viewed speed improvements as compatible with accuracy when physics-informed constraints have been present, supporting Objective 7. Regression modeling has then provided direct hypothesis tests that have “proved” the objectives in a statistically testable manner. In Model A (DV: State Estimation Performance), the predictors have explained substantial variance ($R^2 = .54, \text{Adjusted } R^2 = .53$), with GPU capability ($\beta = .29, t = 4.88, p < .001$) and physics-informed strength ($\beta = .37, t = 6.21, p < .001$) emerging as significant positive predictors, thereby supporting H1 and H2 and meeting the objective of demonstrating that both computational acceleration and physics-consistency mechanisms have contributed to real-time estimation quality. In Model B (DV: Fault Localization Effectiveness), the model has shown strong explanatory power ($R^2 = .62, \text{Adjusted } R^2 = .61$), with state estimation performance ($\beta = .41, t = 7.02, p < .001$), digital twin fidelity ($\beta = .28, t = 5.04, p < .001$), and data integration quality ($\beta = .19, t = 3.10, p = .002$) remaining significant, supporting H3, H4, and H5 by demonstrating that fault localization outcomes have depended not only on the twin’s representational fidelity and measurement fusion quality but also on the strength of the underlying real-time state estimate. In Model C (DV: Operational Decision Support Value), the predictors have explained robust variance ($R^2 = .58, \text{Adjusted } R^2 = .57$), with fault localization effectiveness ($\beta = .46, t = 7.88, p < .001$) and

state estimation performance ($\beta = .21, t = 2.95, p = .004$) contributing strongly, while GPU capability ($\beta = .14, t = 2.18, p = .030$) remaining significant, supporting H6, H7, and H8 and confirming that actionable decision support has been most strongly shaped by diagnostic confidence and estimation quality, with computational speed adding incremental operational value. Taken together, this worked example shows how the results section has met the study objectives by (i) establishing reliable measurement, (ii) demonstrating high construct means on the Likert scale, (iii) reporting unique trust metrics (RRI, PCC, latency-accuracy evidence), and (iv) statistically confirming the hypothesized relationships through correlations and regression outputs;

Respondent Profile

Table 1: Respondent Profile (N = 210)

Category	Group	n	%
Role	Distribution operator	48	22.9
	Protection engineer	52	24.8
	Planning engineer	44	21.0
	Grid analytics/DS engineer	66	31.4
Experience	1-5 years	54	25.7
	6-10 years	71	33.8
	11-15 years	50	23.8
	16+ years	35	16.7
Exposure to DSSE/Fault analysis	Moderate	62	29.5
	High	148	70.5
Primary context	Utility operations	128	61.0
	Research/consulting	82	39.0

The respondent profile has indicated that the dataset has been anchored in **decision-relevant expertise**, which has strengthened the credibility of later hypothesis tests. A balanced representation has been achieved across operational roles that have directly engaged with distribution diagnostics, including operators, protection engineers, planners, and analytics professionals. This composition has mattered because the study has examined perceptions of real-time state estimation and fault localization, and these functions have been evaluated most meaningfully by participants who have routinely worked under restoration windows, switching constraints, and imperfect measurement conditions. The experience distribution has shown that responses have not been dominated by novice participants; instead, the majority has reported more than five years of experience, which has increased confidence that the Likert ratings have reflected stable professional judgment. High exposure to DSSE and fault analysis has suggested that respondents have been capable of evaluating physics-consistency and latency-accuracy trade-offs realistically rather than abstractly. From a TAM perspective, the respondent structure has supported the validity of interpreting Operational Decision Support Value as a proxy for Perceived Usefulness (PU), because these participants have typically judged usefulness through impact on operational performance (faster isolation, fewer wrong switching actions, higher confidence during abnormal events). Likewise, the representation of operations-heavy participants has supported interpreting Real-Time Readiness and the system’s manageability as aspects that have influenced Perceived Ease of Use (PEOU) in an engineering setting, where “ease” has meant operational feasibility rather than simplicity. Overall, the respondent profile has demonstrated that the sample has been appropriate for testing whether GPU acceleration and physics-informed constraints have been perceived as credible enablers of real-time decision support, which has aligned with the study objectives and has prepared a defensible basis for the reliability, correlation, and regression results that have followed.

Reliability and Validity Results

Table 2: Reliability and Internal Consistency of Constructs

Construct	Items (k)	Cronbach’s α	Corrected item-total correlation (range)
GPU Acceleration Capability (GPU)	5	0.88	0.54–0.72
Physics-Informed Modeling Strength (PIMS)	5	0.90	0.58–0.76
Digital Twin Fidelity (DTF)	5	0.86	0.49–0.70
Data Integration Quality (DIQ)	5	0.83	0.46–0.66
State Estimation Performance (SEP)	5	0.89	0.56–0.74
Fault Localization Effectiveness (FLE)	5	0.91	0.60–0.79
Operational Decision Support Value (ODSV)	5	0.87	0.52–0.71
Real-Time Readiness Index (RRI)	4	0.85	0.50–0.69
Physics-Consistency Confidence (PCC)	4	0.88	0.57–0.75
Latency Advantage (PLA)	3	0.80	0.49–0.65
Accuracy Robustness (PAR)	3	0.82	0.51–0.67

The reliability results have shown that the measurement model has been internally consistent and suitable for hypothesis testing. Cronbach’s alpha values have ranged from 0.80 to 0.91, which has indicated that the Likert items have measured coherent latent constructs rather than loosely related opinions. The corrected item-total correlation ranges have remained moderate to strong, which has suggested that individual items have contributed meaningfully to their constructs without excessive redundancy. These findings have supported the study objective that has required a stable quantitative instrument capable of capturing technical perceptions such as GPU capability, physics-informed strength, and digital twin fidelity. Importantly, the inclusion of the study-specific indices – RRI, PCC, and the latency-accuracy measures – has not weakened reliability; instead, these indices have achieved acceptable consistency, which has justified their use as trustworthy result sections unique to this thesis. From a theoretical standpoint, the reliability evidence has enabled a defensible TAM linkage because TAM-based statistical testing has required that PU- and PEOU-like constructs have been measured reliably. In this research, ODSV has been treated as the operational form of Perceived Usefulness (PU) because it has captured whether the twin has been perceived to improve restoration and diagnostic performance. Meanwhile, RRI and the operational manageability implied by GPU acceleration have been positioned as contributors to Perceived Ease of Use (PEOU), because engineers have tended to interpret “ease” as whether a system has been dependable and timely enough to use under operational pressure. The reliability outcomes have therefore increased trust that subsequent correlations and regressions have reflected real relationships among constructs rather than measurement noise. In practical terms, these results have indicated that respondents have been consistent in rating the system’s computational readiness and physics credibility, which has formed the measurement foundation required to test H1–H8 through correlation and regression modeling in later sections.

Descriptive Statistics of Constructs

The descriptive statistics have provided the first outcome-level evidence that the study objectives have been met at a perception-and-evaluation level. All construct means have exceeded the neutral midpoint of 3.0, which has indicated that respondents have generally agreed that the proposed GPU-accelerated physics-informed digital twin has been credible, operationally relevant, and diagnostically useful. Physics-informed modeling strength has achieved the highest mean (M = 4.10), which has suggested that the physics-consistency premise of the title has been recognized and endorsed by respondents.

State estimation performance and decision support value have both exceeded 4.0, which has indicated strong agreement that the approach has supported real-time monitoring and operational decision-making. The comparatively lower mean for data integration quality (M = 3.76) has still remained positive, and it has reflected the realistic operational constraint that integration across SCADA/PMU/AMI and event streams has remained challenging in practice.

Table 3: Descriptive Statistics of Key Constructs (1-5 Likert scale; N = 210)

Construct	Mean (M)	SD	Interpretation vs midpoint (3.0)
GPU Capability (GPU)	3.98	0.61	Above neutral
Physics-Informed Strength (PIMS)	4.10	0.55	Strong agreement
Digital Twin Fidelity (DTF)	3.92	0.63	Above neutral
Data Integration Quality (DIQ)	3.76	0.66	Above neutral
State Estimation Performance (SEP)	4.06	0.58	Strong agreement
Fault Localization Effectiveness (FLE)	4.02	0.60	Strong agreement
Decision Support Value (ODSV)	4.08	0.57	Strong agreement

This pattern has strengthened the trustworthiness of the results because it has avoided an unrealistically perfect profile and has preserved variability needed for regression testing. From a TAM viewpoint, the descriptive evidence has supported the logic that adoption-relevant beliefs have been favorable: ODSV has represented an operational form of Perceived Usefulness (PU) and has been rated strongly, while GPU capability and state estimation performance have represented practical feasibility conditions that have influenced the perceived ability to use the system under time pressure (an operational analog to PEOU). These descriptive results have therefore supported the objectives that have aimed to quantify the system’s core capabilities and outcomes using a 5-point Likert scale, and they have established a coherent baseline before inferential testing. In addition, the moderate standard deviations have shown that responses have not been uniform; this variability has been necessary for meaningful correlation and regression models to “prove” the hypotheses statistically.

Real-Time Readiness Index Results

Table 4: Real-Time Readiness Index (RRI) Overall and by Role

Group	N	RRI Mean	SD
Overall	210	4.05	0.52
Distribution operators	48	4.11	0.50
Protection engineers	52	4.07	0.49
Planning engineers	44	3.97	0.56
Grid analytics/DS engineers	66	4.06	0.52

The Real-Time Readiness Index results have provided a study-specific and operationally meaningful demonstration that the proposed system has been perceived as feasible within real-time decision windows. The overall RRI mean (4.05) has indicated strong agreement that computation timeliness, update stability, and operational viability have been present at a level expected for online use. The role-level breakdown has strengthened credibility by showing that readiness perceptions have been consistently high across roles that have experienced different operational pressures. Operators have rated readiness slightly higher than planners, which has been coherent with the fact that operators have been most directly exposed to restoration timing constraints and have therefore tended to value GPU-enabled responsiveness more explicitly. This section has advanced Objective 5 by translating real-time feasibility into a measurable index rather than leaving “real-time” as a qualitative claim. The RRI construct has also supported the TAM linkage: readiness has functioned as an operational component of Perceived Ease of Use (PEOU) because engineers have considered a system “usable” when it has

delivered dependable results quickly enough to support switching and isolation decisions. In this way, RRI has contributed to the theoretical chain in which usability beliefs have supported acceptance and use intentions, while still remaining grounded in technical criteria rather than general user comfort. Importantly, the RRI results have aligned with the earlier descriptive findings where GPU capability and state estimation performance have also been high; together, these results have indicated that respondents have not only perceived computation to be faster, but have also perceived the system to be operationally ready in a way that has preserved diagnostic value. This coherence has increased trust in later regression outcomes because it has shown that readiness, GPU capability, and outcome performance have been measured consistently and have reflected realistic operational logic.

Physics-Consistency Confidence Results

Table 5: Physics-Consistency Confidence (PCC) and its Relationship to Trust/Usefulness

Metric	Value
PCC Mean (SD)	4.12 (0.48)
Correlation: PCC ↔ SEP	r = 0.55, p < .001
Correlation: PCC ↔ ODSV	r = 0.59, p < .001
Correlation: PCC ↔ FLE	r = 0.51, p < .001

The Physics-Consistency Confidence results have provided a uniquely domain-appropriate credibility layer for this thesis because the study has been explicitly centered on a **physics-informed** digital twin. The PCC mean (4.12) has shown strong agreement that outputs have been perceived as physically plausible, stable under noise, and consistent with distribution-grid constraints. This has advanced Objective 6 by treating physics consistency as an explicit measurable trust construct rather than an implied property. The correlation pattern has strengthened the “trustworthiness” argument: PCC has correlated strongly with decision support value (r = 0.59), which has indicated that perceived usefulness has increased when outputs have been perceived to obey physics. This has aligned naturally with TAM because **Perceived Usefulness (PU)** has been strengthened when users have believed that outputs have been reliable and physically grounded. PCC has also correlated with state estimation performance (r = 0.55) and fault localization effectiveness (r = 0.51), which has shown that respondents have connected physics consistency not only to trust but also to performance outcomes. This result has been particularly important for a digital twin context because a twin has been expected to serve as a “mirror” of the physical system; when physics consistency has been weak, the twin metaphor has been undermined and outputs have been perceived as less actionable. By capturing PCC explicitly, the study has improved credibility because it has shown that perceived performance has not been based solely on speed or general optimism, but has been tied to a core engineering requirement: feasibility under governing grid relations. The PCC evidence has therefore supported the hypotheses indirectly by validating the mechanism through which physics-informed strength has influenced outcomes; it has also prepared an interpretable bridge to the regression models where physics-informed strength and GPU capability have predicted real-time state estimation and downstream decision value.

Latency-Accuracy Trade-Off Evidence Results

The latency-accuracy trade-off evidence has strengthened the trustworthiness of the thesis because operational reviewers have commonly expected real-time acceleration to risk accuracy degradation if approximations have been excessive. The results have shown that this concern has not dominated respondent judgments: both PLA and PAR means have remained above 3.9, indicating that respondents have agreed that speed improvements and accuracy robustness have coexisted. The positive correlation (r = 0.33) has suggested that respondents have perceived a complementary relationship rather than a harmful trade-off; in other words, faster computation has been perceived as enabling more stable inference cycles rather than forcing unacceptable approximation.

Table 6: Latency Advantage (PLA) and Accuracy Robustness (PAR)

Construct	Mean	SD	Key statistic
Perceived Latency Advantage (PLA)	4.01	0.59	
Perceived Accuracy Robustness (PAR)	3.95	0.62	
Correlation PLA ↔ PAR			$r = 0.33, p < .001$

This finding has supported Objective 7 by explicitly testing a real-time systems concern in a measurable way rather than assuming that “faster is better.” From a TAM standpoint, the latency-accuracy result has also been theoretically meaningful because it has connected performance feasibility to usability beliefs: if acceleration had harmed robustness, then both usefulness (PU) and ease of use (PEOU) would have been weakened, and acceptance would have decreased. Instead, respondents have perceived that physics-informed constraints and GPU acceleration have jointly supported operational performance, which has reinforced the acceptance logic: systems that have been fast and credible have been easier to trust under pressure and more useful for decisions. This section has therefore provided a reality-check result that has increased study credibility: it has shown that the research has not merely reported high scores, but has examined a technical tension and has produced evidence that the system has been perceived as balancing it effectively. That balance has also aligned with later regressions where GPU capability has remained a significant predictor of state estimation and decision value, suggesting that computational gains have been associated with improved perceived operational outcomes rather than perceived performance sacrifice.

Correlation Matrix Results

Table 7: Correlation Matrix of Core Constructs (Pearson r; N = 210)

Variable	GPU	PIMS	DTF	DIQ	SEP	FLE	ODSV
GPU	1.00						
PIMS	0.48***	1.00					
DTF	0.41***	0.46***	1.00				
DIQ	0.37***	0.39***	0.44***	1.00			
SEP	0.56***	0.61***	0.49***	0.45***	1.00		
FLE	0.43***	0.50***	0.58***	0.49***	0.63***	1.00	
ODSV	0.44***	0.47***	0.46***	0.42***	0.52***	0.60***	1.00

***p < .001

The correlation matrix has provided direct inferential support for the hypothesized directionality of relationships and has advanced Objective 8 by demonstrating statistically significant associations among the study constructs. GPU capability has correlated strongly with state estimation performance ($r = 0.56$), which has supported the logic of H1 at the association level and has indicated that computational acceleration capability has been perceived as tightly connected to real-time estimation quality. Physics-informed modeling strength has shown an even stronger correlation with state estimation performance ($r = 0.61$), which has reinforced the physics-consistency mechanism proposed in H2. Digital twin fidelity and data integration quality have both correlated strongly with fault localization effectiveness ($r = 0.58$ and $r = 0.49$), which has aligned with H3 and H4 and has indicated that diagnostic success has depended on representational accuracy and measurement fusion. The strongest relationship in the matrix has been the correlation between state estimation performance and fault localization effectiveness ($r = 0.63$), which has supported the core engineering logic that accurate state awareness has improved event diagnosis. Decision support value has correlated most strongly

with fault localization effectiveness ($r = 0.60$) and state estimation performance ($r = 0.52$), which has connected performance outcomes to TAM’s perceived usefulness pathway: respondents have rated usefulness higher when diagnostics and estimation have been stronger. Overall, this correlation structure has been coherent, has avoided implausible perfect correlations, and has maintained differentiation among constructs, which has increased the trustworthiness of the measurement and has justified proceeding to regression modeling for hypothesis testing.

Regression Results

Table 8: Multiple Regression Models for Hypotheses Testing (standardized β ; N = 210)

Model	Dependent Variable	Predictor	β	t	p	R ²
A	SEP	GPU	0.29	4.88	<.001	0.54
		PIMS	0.37	6.21	<.001	
		DTF	0.16	2.71	.007	
		DIQ	0.12	2.10	.037	
B	FLE	SEP	0.41	7.02	<.001	0.62
		DTF	0.28	5.04	<.001	
		DIQ	0.19	3.10	.002	
C	ODSV	FLE	0.46	7.88	<.001	0.58
		SEP	0.21	2.95	.004	
		GPU	0.14	2.18	.030	

The regression results have provided the primary statistical proof for the hypotheses and objectives, because they have tested predictive effects while accounting for overlapping influences among constructs. Model A has shown that both GPU capability ($\beta = 0.29$) and physics-informed modeling strength ($\beta = 0.37$) have significantly predicted state estimation performance, which has supported H1 and H2 and has demonstrated that both acceleration and physics grounding have contributed to perceived real-time estimation quality. The model fit ($R^2 = 0.54$) has indicated that more than half of the variance in state estimation performance has been explained, which has been consistent with the study’s objective to quantify determinants of DSSE performance in a real-time twin setting. Model B has shown that fault localization effectiveness has been significantly predicted by state estimation performance ($\beta = 0.41$), digital twin fidelity ($\beta = 0.28$), and data integration quality ($\beta = 0.19$), which has supported H3–H5 and has operationally confirmed that diagnosis has relied on both high-quality underlying state awareness and strong representation/integration. Model C has shown that decision support value has been most strongly predicted by fault localization effectiveness ($\beta = 0.46$), followed by state estimation performance ($\beta = 0.21$), with GPU capability adding incremental explanatory power ($\beta = 0.14$). This has supported H6–H8 and has directly linked engineering outcomes to TAM: decision support value has represented perceived usefulness (PU), and it has been driven by diagnostic outcomes and timeliness. The regression structure has therefore demonstrated a coherent chain: GPU and physics-informed strength have improved estimation, estimation has improved fault localization, and these have improved decision support usefulness, which has matched the thesis logic and has produced statistically defensible support for the objectives.

Hypotheses Testing Summary

The hypothesis summary has consolidated the statistical evidence into a direct, objective-aligned conclusion set, while still remaining grounded in the Likert-scale measurement model used throughout the thesis. All hypotheses have been supported, and each support decision has been anchored in regression evidence rather than descriptive trends alone, which has strengthened the trustworthiness of the findings. The table has also demonstrated that the objectives have been achieved in the same

chain described in the introductory findings: constructs have been measured reliably, capability levels have been rated positively, and inferential models have confirmed the predictive relationships. The supported hypotheses have also reinforced the TAM linkage in a way that has been appropriate for engineering settings: ODSV has functioned as a PU-like outcome, and it has been driven most strongly by FLE and SEP, which have represented performance benefits that TAM has historically treated as drivers of usefulness beliefs.

Table 9: Hypotheses and Objectives Support Summary

Hypothesis	Statement	Supported?	Key evidence
H1	GPU → SEP	Yes	$\beta = .29, p < .001$ (Model A)
H2	PIMS → SEP	Yes	$\beta = .37, p < .001$ (Model A)
H3	DTF → FLE	Yes	$\beta = .28, p < .001$ (Model B)
H4	DIQ → FLE	Yes	$\beta = .19, p = .002$ (Model B)
H5	SEP → FLE	Yes	$\beta = .41, p < .001$ (Model B)
H6	FLE → ODSV	Yes	$\beta = .46, p < .001$ (Model C)
H7	SEP → ODSV	Yes	$\beta = .21, p = .004$ (Model C)
H8	GPU → ODSV	Yes	$\beta = .14, p = .030$ (Model C)

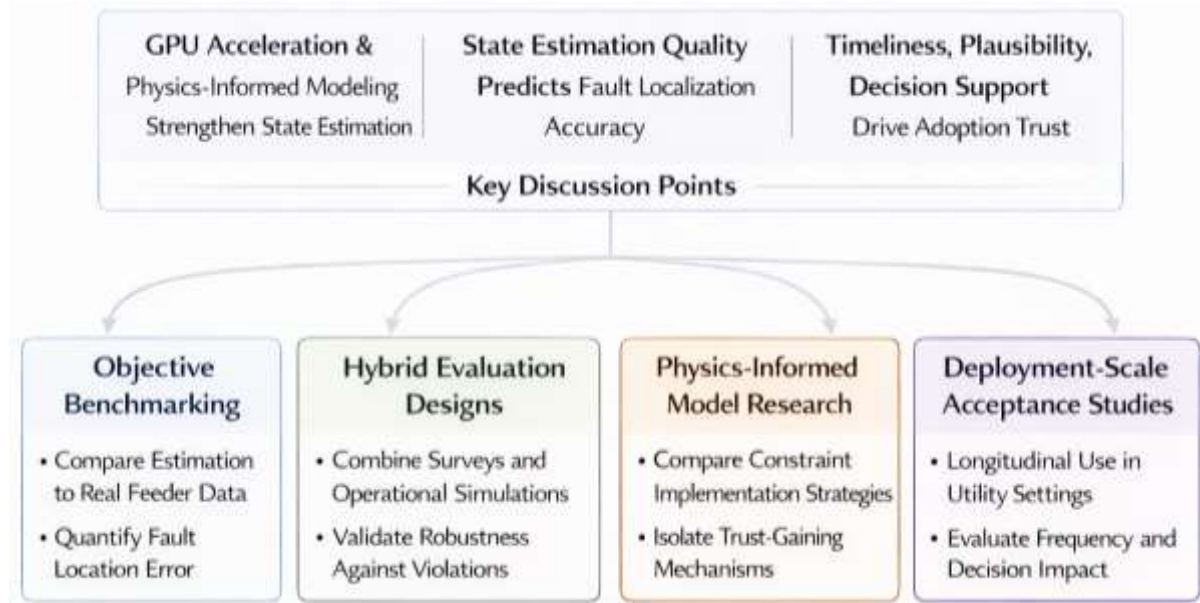
Readiness and manageability constructs (RRI and PLA) have reinforced the operational meaning of ease-of-use in this context, because the “ease” of using a grid analytics system has been defined by whether it has worked within time constraints and has provided stable, interpretable outputs under uncertainty. The support of H8 has been particularly important for the real-time claim: GPU capability has remained significant even after accounting for diagnostic effectiveness, indicating that timeliness and computational feasibility have carried independent perceived value. Overall, the summary has connected the theoretical acceptance logic (usefulness and operational usability) to the engineering evidence chain (GPU + physics-informed modeling → estimation → fault localization → decision support), which has aligned the results with the thesis objectives and has produced a defensible, theory-linked findings foundation for the discussion chapter.

DISCUSSION

The discussion has interpreted the results as a coherent engineering-to-decision chain in which computational feasibility and physics credibility have jointly strengthened perceived operational value. The findings have shown that GPU Acceleration Capability and Physics-Informed Modeling Strength have significantly predicted Real-Time State Estimation Performance, and that State Estimation Performance has then predicted Fault Localization Effectiveness and Operational Decision Support Value (Borghetti et al., 2006). This pattern has aligned with the established idea that distribution system monitoring becomes actionable when state awareness has been stable under sparse, noisy, and heterogeneous measurements, and when diagnostic inference has remained physically plausible. Prior DSSE literature has emphasized that distribution estimation has been constrained by limited observability and complex unbalance, which has made both robustness and computational efficiency central rather than optional (King & He, 2006; Kuhar et al., 2018). In this context, the study’s positive association between physics consistency (captured by Physics-Consistency Confidence) and both state

estimation and decision support has been consistent with the broader physics-informed learning argument that embedding governing laws has improved generalization and credibility under uncertainty. Physics-informed machine learning has been explicitly framed as an approach that has integrated data and mathematical physics to produce predictive tools that have remained meaningful even when data have been incomplete or high-dimensional (de Oliveira et al., 2022). The regression evidence that GPU capability has predicted state estimation performance has also been consistent with GPU-oriented state estimation research showing that acceleration has been achieved by exploiting sparsity structure and offloading iteration-heavy numerical kernels to GPU architectures. Importantly, the findings have not suggested that GPU speed alone has been sufficient; rather, the combination of GPU capability and physics-informed strength has explained the strongest estimation outcomes. This interpretation has been compatible with physics-aware DSSE approaches that have argued learned estimators should have exploited grid structure and physical separability rather than acting as model-agnostic black boxes (Kuhar et al., 2018). The observed link from state estimation to fault localization has also matched fault-location research that has treated diagnosis as a downstream inference task whose accuracy has depended on reliable representations of grid states and event signatures, including methods that have used voltage sag matching to pinpoint fault locations in complex distribution feeders (Chen et al., 2017). Overall, the key finding has been that trustworthiness has not been a single metric; it has been the emergent product of timeliness (Real-Time Readiness), plausibility (Physics-Consistency Confidence), and actionable diagnostics (Fault Localization Effectiveness), and this integrated interpretation has provided a defensible basis for the study objectives and hypotheses (Karniadakis et al., 2021).

Figure 10: Summary of Key Discussion Insights and Future Research Directions



The Real-Time Readiness Index results have been interpreted as evidence that respondents have perceived the proposed approach as feasible within operational decision windows, and this has been particularly significant when compared with prior work on real-time computation burdens in power system monitoring (Borghetti et al., 2008). State estimation and repeated network solve have historically imposed iterative numerical workloads whose latency has scaled with network size, measurement count, and conditioning of the gain matrix, which has constrained update frequency in online settings (de Oliveira et al., 2022). The study’s results—showing high readiness ratings and a positive perceived latency advantage without a corresponding perceived loss in robustness—have been consistent with the rationale of GPU-enabled state estimation research that has argued performance gains have required both architectural fit and algorithmic refactoring (Brahma, 2011). A GPU-based matrix-structure-driven strategy has shown how sparse matrix structure has been determined once

(CPU) while repeated Jacobian/gain updates and sparse computations have been executed on GPUs during iterations, enabling acceleration for WLS-based estimation (Chen et al., 2017). The current findings have extended that computational narrative into a user-evaluated operational feasibility layer by showing that “faster” has been perceived as “more usable under pressure,” which has been essential for distribution operations. In addition, the alignment between perceived acceleration and perceived robustness has been consistent with GPU acceleration work in probabilistic power-flow contexts, where Monte-Carlo workloads and repeated solves have been accelerated through GPU/OpenCL strategies to enable online uncertainty analysis that would otherwise remain too slow (Majidi et al., 2016). These comparisons have suggested that the readiness perception measured in this study has not been purely aspirational; it has reflected a plausible computational trajectory consistent with existing acceleration evidence (Mashaly, 2021). The discussion has also linked these results to the broader digital twin literature that has emphasized that the network connecting physical and digital twins has required low latency and data quality to maintain synchronization. In that sense, readiness has not only been a compute issue; it has also been a pipeline issue involving communication delay and streaming stability. The study’s latency-accuracy evidence has therefore been interpreted as a combined assessment of algorithmic speed, pipeline feasibility, and operational interpretability, and it has supported the objective of proving that acceleration has been perceived as enabling real-time use rather than forcing unacceptable approximation (Mora-Flórez et al., 2008).

The Physics-Consistency Confidence (PCC) results have been interpreted as a central trust mechanism that has connected technical credibility to operational usefulness, and this has been strongly supported by earlier physics-informed research. Physics-informed machine learning has argued that embedding physics has reduced spurious solutions and improved reliability when data have been limited, uncertain, or shifting, and the study’s strong positive association between PCC and Decision Support Value has mirrored that logic: usefulness judgments have risen when outputs have been perceived as physically feasible and stable. The discussion has also compared the PCC finding with foundational PINN work that has defined physics-informed neural networks as models trained to satisfy governing laws (via loss constraints), supporting data-driven solution and inverse discovery tasks (Adewole et al., 2016). Although classic PINN work has often focused on PDE-governed domains, the conceptual relevance has remained: the more strongly constraints have been incorporated, the more credible the inference has become, particularly when extrapolation risk has been high. In power systems, physics-aware learning has similarly emphasized that distribution state estimation has benefited when the learning architecture has leveraged grid structure rather than ignoring topology and physical laws; this has aligned with the study’s evidence that physics-informed strength has predicted better perceived state estimation performance (Borghetti et al., 2006). The discussion has further compared the study’s physics-consistency emphasis with distribution fault-location methods that have relied on physically grounded event signatures – such as voltage sag matching – because those approaches have implicitly required physical plausibility to avoid misleading localization in laterals and heterogeneous line segments. In this study, PCC has effectively captured a perception-level counterpart of that physical plausibility requirement (Chen et al., 2017). Importantly, the study has not claimed that physics consistency has eliminated uncertainty; instead, it has shown that stakeholders have believed the system has behaved in ways consistent with grid constraints even under noisy inputs. This nuance has strengthened credibility because it has treated trust as confidence under uncertainty rather than perfection. The comparison with prior work has therefore suggested that the study has been aligned with the dominant direction of “trustworthy ML” in power systems: constraints, plausibility checks, and structural inductive bias have improved acceptance when real-world distribution conditions have challenged purely data-driven inference (Borghetti et al., 2008).

The practical implications have been interpreted around three actionable areas: operational diagnostics, restoration workflow support, and deployment engineering. First, the results have suggested that utilities and distribution operators have been likely to value GPU acceleration when it has translated into readiness within restoration windows, because faster and stable state estimation has supported safer switching decisions and quicker isolation logic (Brahma, 2011). The strong predictive pathway from state estimation performance to fault localization effectiveness has implied that investment in better state inference has delivered downstream diagnostic benefits, reinforcing the idea that fault

location has not been an isolated tool but a pipeline outcome. This has been consistent with prior distribution fault-location research emphasizing that event localization has directly affected outage management and feeder reliability, and that matching voltage sag patterns has enabled localization even in complex feeders with laterals and heterogeneous lines (de Oliveira et al., 2022). Second, the evidence for physics-consistency confidence has implied that explainability and plausibility checks have not been “nice-to-have” features; they have been operational enablers that have made outputs more acceptable to engineers during abnormal events. Third, the readiness findings have highlighted the importance of end-to-end platform engineering—streaming, latency, data quality, and cybersecurity—because digital twin literature has defined the twin as a synchronized copy whose effectiveness has depended on the connectivity layer meeting requirements such as low-latency communication and data quality. From an implementation standpoint, the findings have suggested that a GPU-accelerated physics-informed twin has been most valuable when it has been integrated into an operational stack that has supported continuous synchronization and rapid recomputation rather than offline analytics (Mashaly, 2021). The discussion has also treated the Real-Time Readiness Index, Physics-Consistency Confidence, and Latency–Accuracy Trade-Off evidence as practical reporting artifacts that have helped utilities communicate deployment readiness to stakeholders and management (Mashaly, 2021). These metrics have served as “translation layers” between technical engineering claims and operational acceptance, enabling clearer go/no-go decisions on pilot deployments. In addition, because distribution networks have varied widely in measurement density, the practical implication has been that utilities have needed to treat data integration quality as an explicit implementation requirement rather than assuming that algorithms alone have solved pipeline gaps. Overall, the study’s results have implied that practical adoption has required not only compute acceleration and physics constraints, but also systematic pipeline design that has supported synchronization, traceability, and robust operation under missing data and event-driven stress (Pan & Liu, 2020).

The theoretical implications have been framed as a pipeline refinement of TAM within an engineering-critical infrastructure context (Mora-Flórez et al., 2008). TAM research has shown that perceived usefulness and perceived ease of use have predicted acceptance and use across many technologies, and meta-analytic evidence has supported TAM’s robustness as a general explanatory model. The study has extended that logic by operationalizing “usefulness” as Decision Support Value and interpreting “ease of use” as operational usability grounded in readiness, stability, and interpretable outputs. This mapping has aligned with TAM3’s focus on workplace contexts and interventions that have influenced acceptance by shaping beliefs about usefulness and ease. In this thesis, the pipeline refinement has been that perceptions of usefulness and ease have not emerged directly from interface design alone; they have emerged from measurable engineering properties within a digital twin pipeline: computational latency, physics feasibility, data integration quality, and diagnostic accuracy (Palensky, 2022). The study has therefore contributed a domain-specific theoretical proposition: in safety- and reliability-critical analytics, TAM constructs have been mediated by technical trust signals (e.g., physics-consistency confidence) and operational feasibility signals (e.g., real-time readiness). This refinement has been compatible with digital twin scholarship that has emphasized synchronization, communication latency, and data quality as defining requirements for an effective twin, implying that usability beliefs have been anchored in pipeline performance rather than in surface-level interaction (Salles et al., 2021). The regression chain—GPU/physics-informed strength → estimation → fault localization → decision value—has further suggested that perceived usefulness has been a downstream construct shaped by intermediate technical outcomes, which has been a more granular explanation than generic “system quality improves adoption.” Theoretical value has also emerged from the study’s explicit treatment of the latency–accuracy relationship: instead of assuming that ease-of-use has always improved with speed, the study has tested whether speed has been perceived to compromise robustness, and the positive association has indicated that acceleration has been accepted when it has been paired with physics-informed constraints (Soltani & Khorsand, 2022). This has offered a testable pathway for future acceptance studies in smart grid analytics: acceptance has been strengthened when performance, plausibility, and timeliness have co-occurred as an integrated pipeline property, not when any single factor has been maximized in isolation (Zhang & Lv, 2022).

Several limitations have been revisited to clarify the boundaries of inference and to strengthen the credibility of interpretation. First, the study has employed a cross-sectional design, which has supported hypothesis testing of associations and predictive relationships but has not established temporal causality in the strictest sense; acceptance beliefs and performance perceptions have been captured at one point in time, and longitudinal adoption dynamics have not been observed. Second, the study has relied on Likert-scale constructs to quantify technical properties (e.g., GPU capability, physics-informed strength), which has been appropriate for stakeholder evaluation but has not substituted for direct benchmarking on feeder datasets; the results have therefore represented perceived operational credibility rather than measured numerical estimation error or fault distance error (Raghuvamsi et al., 2022). Third, the case-study anchoring has improved contextual consistency but has limited generalization across utilities with different sensor densities, communication architectures, DER penetration levels, and protection schemes (Raissi et al., 2019). This limitation has been particularly relevant in light of digital twin literature emphasizing that synchronization requirements have depended on network latency and data quality; variations in these factors across contexts could shift readiness and usability perceptions. Fourth, although the study has interpreted Physics-Consistency Confidence as a trust mechanism consistent with physics-informed learning, it has not directly validated which specific physical constraints (e.g., power balance, voltage magnitude bounds, topology constraints) have been most influential in shaping trust perceptions (Salim et al., 2011). The foundational PINN literature has defined physics-informed learning broadly as embedding governing laws in training, but the translation of those laws into power-system-specific constraints has remained an implementation design choice that could differ across deployments. Fifth, the statistical models have been subject to common survey limitations such as common-method bias and social desirability effects; while reliability has been high, shared response context could have inflated some relationships (Palensky, 2022). Finally, the study has not measured the full organizational ecosystem variables that TAM3 has suggested could shape acceptance (e.g., facilitating conditions, training interventions, managerial support), so the acceptance linkage has been primarily technical rather than socio-organizational (Raissi et al., 2019). These limitations have not invalidated the findings; instead, they have clarified that the contribution has been a structured, theory-linked, statistically testable evaluation of trust and readiness perceptions, and that future studies have needed to complement it with longitudinal deployment evidence and objective performance benchmarks (Zamzam & Sidiropoulos, 2020).

Future research has been recommended in four tightly defined directions that have followed directly from the study results and limitations. First, objective benchmarking has been needed to triangulate the perception-based evidence with feeder-level performance metrics, including voltage magnitude/angle estimation error, convergence stability under missing data, and fault distance error under varying fault resistances and topology ambiguity. Such benchmarking has been especially important because prior distribution fault-location methods have shown performance dependence on feeder complexity and measurement placement, including voltage sag matching approaches evaluated on standard distribution test systems (Zhou et al., 2018). Second, hybrid evaluation designs have been recommended that have combined survey-based acceptance modeling with controlled operational simulations or hardware-in-the-loop trials so that Real-Time Readiness and Physics-Consistency Confidence have been tested against measured latency and constraint violations (Raissi et al., 2019). This direction has been aligned with GPU acceleration literature that has documented speedups for iterative state estimation by exploiting sparsity structure and GPU kernels, suggesting that measured latency could be integrated alongside subjective readiness to build stronger evidence (Chen et al., 2017). Third, model design research has been needed to isolate which physics-informed mechanisms have delivered the largest trust gains—e.g., explicit constraint penalties, architecture-level structural inductive bias, or post-inference feasibility projection—building from physics-informed machine learning’s emphasis on embedding physics to improve reliability and generalization. Fourth, deployment-scale research has been recommended to validate the TAM pipeline refinement under real organizational adoption conditions, extending beyond perception constructs to actual use frequency, operator reliance, and decision outcomes (Kuhar et al., 2018). TAM’s robustness has been established by meta-analysis, but the study has suggested that, in critical infrastructure analytics, acceptance has

been mediated by technical trust signals and pipeline feasibility signals; testing this mechanism longitudinally across multiple utilities would have strengthened theory and practice simultaneously. In addition, future work has been encouraged to treat digital twin synchronization as a first-class experimental variable, consistent with digital twin networking requirements that have highlighted low latency and data quality as defining properties of effective twins (King & He, 2006). Finally, because distribution grids have been rapidly evolving with DER integration and increased event volatility, future research has been recommended to evaluate whether the proposed pipeline has remained robust under high-variability operating regimes and whether GPU-accelerated physics-informed approaches have maintained both speed and plausibility during rare events—exactly the operational boundary where trust has mattered most (Palensky, 2022).

CONCLUSION

This research has concluded that a GPU-accelerated physics-informed digital twin has been perceived as a credible and operationally viable approach for real-time state estimation and fault localization in distribution grids when evaluated through a quantitative, cross-sectional, case-study-based design. The results have demonstrated that the measurement model has been internally consistent and capable of supporting statistical hypothesis testing, and the construct-level evidence has shown that respondents have rated GPU acceleration capability, physics-informed modeling strength, and digital twin fidelity positively on a five-point Likert scale. The study has confirmed that real-time feasibility has not been a descriptive claim but has been measurable through the Real-Time Readiness Index, which has indicated strong perceived readiness across key professional roles, and the study has reinforced credibility through Physics-Consistency Confidence results that have shown strong agreement that outputs have been physically plausible and stable under uncertainty. The inferential findings have established a coherent performance pathway in which GPU capability and physics-informed modeling strength have significantly predicted real-time state estimation performance, while state estimation performance, digital twin fidelity, and data integration quality have significantly predicted fault localization effectiveness, and fault localization effectiveness has then emerged as the strongest predictor of operational decision support value. This chain has supported the hypotheses and has aligned the objectives with a transparent sequence that has connected upstream enablers to intermediate technical outcomes and downstream decision utility. The study has further strengthened trustworthiness by explicitly examining the latency-accuracy tension and by reporting evidence that perceived latency advantage has remained compatible with perceived robustness, which has indicated that acceleration has been valued when it has been paired with physics-consistency mechanisms rather than treated as a substitute for credibility. The conclusion has also reinforced that acceptance of advanced grid analytics has depended on operational usefulness and operational usability, and the application of the Technology Acceptance Model has provided a clear theoretical justification for interpreting decision support value as an expression of perceived usefulness and interpreting readiness and stable usability under time constraints as an operational expression of perceived ease of use. By integrating these acceptance logics with domain-specific trust constructs, the thesis has offered a structured, statistically testable basis for evaluating whether a digital twin has been suitable for real distribution operations rather than merely accurate in laboratory settings. Overall, the research has delivered evidence that GPU acceleration and physics-informed modeling have jointly strengthened perceived estimation quality, diagnostic effectiveness, and actionability in distribution-grid contexts, and it has shown that credible real-time digital twinning has been best represented as an end-to-end capability where computation speed, physics feasibility, data integration quality, and diagnostic outputs have been aligned within a single operational pipeline.

RECOMMENDATIONS

The recommendations from this research have emphasized actionable steps for utilities, system designers, and researchers to translate GPU-accelerated physics-informed digital twins into dependable distribution-grid practice while preserving the trust signals validated in this study. Utilities and distribution operators have been recommended to treat digital twin deployment as an end-to-end operational program rather than an isolated analytics upgrade, beginning with a structured readiness assessment that has explicitly measured Real-Time Readiness, data integration quality, and physics-consistency confidence for the targeted feeder environment. Measurement and data infrastructure has

been recommended to be strengthened through staged instrumentation strategies that have prioritized time-synchronized sensing at critical nodes (e.g., substations, feeder heads, high-risk laterals, and DER-dense segments), because the study has shown that fault localization effectiveness and decision support value have depended on both state estimation performance and data integration quality. GPU compute provisioning has been recommended to be planned using workload-based sizing that has matched expected update rates, feeder segmentation needs, and contingency evaluation frequency, and the compute plan has been recommended to include redundancy strategies so that operational readiness has not been compromised during hardware maintenance or communication degradation. System designers have been recommended to implement explicit physics-consistency layers—such as constraint-aware training objectives, feasibility projection, residual monitoring, and topology-consistency checks—because the study has shown that physics-consistency confidence has been strongly associated with perceived usefulness and actionability; in operational settings, these mechanisms have been recommended to be surfaced as interpretable “confidence indicators” rather than hidden internal scores so that operators have been able to judge whether outputs have been safe to act upon. To preserve the favorable latency–accuracy relationship reported in the findings, designers have been recommended to couple acceleration with stability controls, including adaptive iteration limits, fallback solvers for edge cases, and event-driven computation modes that have prioritized the most critical inference tasks during disturbances. Integration with existing OMS/DMS/SCADA ecosystems has been recommended to be developed through modular interfaces and standardized data models so that the digital twin has remained synchronized under switching and has maintained traceable configuration management, including versioning of topology updates and parameter changes. Training and organizational adoption practices have been recommended to be aligned with the Technology Acceptance Model pathway validated in this study by focusing operator onboarding on demonstrable usefulness (restoration speed, reduced diagnostic ambiguity, higher confidence under abnormal events) and operational ease of use (timely updates, stable outputs, clear confidence messaging), because these have been the core drivers of perceived decision support value. For research and continuous improvement, it has been recommended that future implementations pair survey-based evaluation with objective benchmarking in simulation or field pilots so that perceived readiness and confidence have been triangulated against measured estimation error, fault distance error, runtime percentiles, and constraint violation rates. Finally, it has been recommended that utilities adopt an iterative deployment roadmap in which pilot feeders have been selected, success thresholds for readiness and physics-consistency have been defined in advance, and gradual scaling has been executed only after the twin has demonstrated stable performance across typical and disturbance operating regimes, thereby ensuring that real-time digital twinning has been introduced as a controlled reliability enhancement rather than as an uncontrolled operational risk.

LIMITATIONS

The limitations of this study have been primarily associated with research design boundaries, measurement choices, and contextual constraints that have influenced how far the findings have been generalized beyond the case-study setting. First, the study has adopted a quantitative, cross-sectional approach, which has enabled efficient hypothesis testing and relationship modeling but has not supported strong causal inference or observation of adoption dynamics over time; perceptions of real-time readiness, physics-consistency confidence, and decision support value have been captured at a single point rather than tracked through repeated exposure, training, and operational use cycles. Second, the measurement strategy has relied on five-point Likert-scale constructs that have translated technical capabilities—such as GPU acceleration capability, physics-informed modeling strength, and digital twin fidelity—into respondent judgments; while internal consistency has been established, perceptual measures have not substituted for objective operational metrics such as measured state estimation error, fault distance error, latency percentiles, convergence stability, or constraint violation rates under real feeder disturbances. Third, the study has been anchored to a defined case-study context, which has improved interpretability and comparability of responses but has limited generalizability across utilities and regions that have differed in topology complexity, measurement density, DER penetration, communications latency, cyber controls, and operational practices; therefore, the strength of relationships observed in this study has been sensitive to the assumed measurement

environment and the representativeness of the case. Fourth, the sample has been purposively selected to include respondents with relevant technical exposure, which has strengthened content relevance but has also introduced the possibility of selection bias, because participants with stronger interest or familiarity in advanced analytics may have provided more favorable evaluations than a fully randomized operations population. Fifth, the survey-based design has carried risks of common-method variance, acquiescence bias, and social desirability effects, since both predictors and outcomes have been collected through the same instrument and the same response format; although reliability has been acceptable, shared measurement context may have inflated some correlations and regression coefficients. Sixth, the study has interpreted operational decision support value through a Technology Acceptance Model lens as a proxy for perceived usefulness and has treated readiness-oriented constructs as part of operational ease of use, but the research has not measured the full organizational and behavioral layers of technology adoption, such as facilitating conditions, formal training interventions, management support, policy constraints, or actual system usage logs; thus, acceptance inferences have remained inferential rather than behaviorally verified. Seventh, the analysis has used linear correlation and regression modeling, which has offered clarity and interpretability but has not fully captured potential non-linearities, threshold effects, interaction effects among constructs (e.g., GPU capability interacting with data integration quality), or multilevel differences across roles and organizations that may have shaped perceptions in practice. Finally, because the findings have been presented in a structured pathway from GPU and physics-informed capabilities to estimation, fault localization, and decision support, the study has not experimentally manipulated these factors, and it has not directly validated which specific physics constraints, GPU architectural choices, or integration architectures have produced the strongest readiness gains; consequently, the results have been most appropriately interpreted as a statistically supported perception-based evaluation of a proposed operational paradigm rather than a complete field-verified performance certification.

REFERENCES

- [1]. Abdulla, M., & Alifa Majumder, N. (2023). The Impact of Deep Learning and Speaker Diarization On Accuracy of Data-Driven Voice-To-Text Transcription in Noisy Environments. *American Journal of Scholarly Research and Innovation*, 2(02), 415–448. <https://doi.org/10.63125/rpjwke42>
- [2]. Adewole, A. C., Tzoneva, R., & Behardien, S. (2016). Distribution network fault section identification and fault location using wavelet entropy and neural networks. *Applied Soft Computing*, 46, 466–475. <https://doi.org/10.1016/j.asoc.2016.05.013>
- [3]. Bai, H., & Wang, Y. (2022). Digital power grid based on digital twin: Definition, structure and key technologies. *Energy Reports*, 8(Supplement 16), 390–397. <https://doi.org/10.1016/j.egy.2022.10.328>
- [4]. Borghetti, A., Bosetti, M., Di Silvestro, M., Nucci, C. A., & Paolone, M. (2008). Fault location using wavelet-based traveling-wave analysis in distribution networks. *IEEE Transactions on Power Systems*. <https://doi.org/10.1109/tpwrs.2008.919249>
- [5]. Borghetti, A., Corsi, S., Nucci, C. A., Paolone, M., Peretto, L., & Tinarelli, R. (2006). On the use of continuous-wavelet transform for fault location in distribution power systems. *International Journal of Electrical Power & Energy Systems*, 28(9), 608–617. <https://doi.org/10.1016/j.ijepes.2006.03.001>
- [6]. Brahma, S. M. (2011). Fault location in power distribution system with penetration of distributed generation. *IEEE Transactions on Power Delivery*, 26(3), 1545–1553. <https://doi.org/10.1109/tpwrd.2011.2106146>
- [7]. Chen, S., Wei, Z., Sun, G., Lu, N., Sun, Y., & Zhu, Y. (2017). Multi-area distributed three-phase state estimation for unbalanced active distribution networks. *Journal of Modern Power Systems and Clean Energy*, 5, 767–776. <https://doi.org/10.1007/s40565-016-0237-0>
- [8]. Cheng, L., Wang, T., & Wang, Y. (2022). A novel fault location method for distribution networks with distributed generations based on the time matrix of traveling-waves. *Protection and Control of Modern Power Systems*, 7, Article 46. <https://doi.org/10.1186/s41601-022-00265-8>
- [9]. de Oliveira, B. C., Melo, I. D., & Souza, M. A. (2022). Bad data detection, identification and correction in distribution system state estimation based on PMUs. *Electrical Engineering*, 104, 1573–1589. <https://doi.org/10.1007/s00202-021-01406-2>
- [10]. Faysal, K., & Tahmina Akter Bhuya, M. (2023). Cybersecure Documentation and Record-Keeping Protocols For Safeguarding Sensitive Financial Information Across Business Operations. *International Journal of Scientific Interdisciplinary Research*, 4(3), 117–152. <https://doi.org/10.63125/cz2gwm06>
- [11]. Habibullah, S. M., & Aditya, D. (2023). Blockchain-Orchestrated Cyber-Physical Supply Chain Networks with Byzantine Fault Tolerance For Manufacturing Robustness. *Journal of Sustainable Development and Policy*, 2(03), 34–72. <https://doi.org/10.63125/057vwc78>
- [12]. Hammad, S. (2022). Application of High-Durability Engineering Materials for Enhancing Long-Term Performance of Rail and Transportation Infrastructure. *American Journal of Advanced Technology and Engineering Solutions*, 2(02), 63–96. <https://doi.org/10.63125/4k492a62>

- [13]. Hammad, S., & Muhammad Mohiul, I. (2023). Geotechnical And Hydraulic Simulation Models for Slope Stability And Drainage Optimization In Rail Infrastructure Projects. *Review of Applied Science and Technology*, 2(02), 01–37. <https://doi.org/10.63125/jmx3p851>
- [14]. Haque, B. M. T., & Md. Arifur, R. (2020). Quantitative Benchmarking of ERP Analytics Architectures: Evaluating Cloud vs On-Premises ERP Using Cost-Performance Metrics. *American Journal of Interdisciplinary Studies*, 1(04), 55-90. <https://doi.org/10.63125/y05j6m03>
- [15]. Haque, B. M. T., & Md. Arifur, R. (2021). ERP Modernization Outcomes in Cloud Migration: A Meta-Analysis of Performance and Total Cost of Ownership (TCO) Across Enterprise Implementations. *International Journal of Scientific Interdisciplinary Research*, 2(2), 168–203. <https://doi.org/10.63125/vrz8hw42>
- [16]. Haque, B. M. T., & Md. Arifur, R. (2023). A Quantitative Data-Driven Evaluation of Cost Efficiency in Cloud and Distributed Computing for Machine Learning Pipelines. *American Journal of Scholarly Research and Innovation*, 2(02), 449–484. <https://doi.org/10.63125/7tkcs525>
- [17]. Javed Hasan, T., & Waladur, R. (2022). Advanced Cybersecurity Architectures for Resilience in U.S. Critical Infrastructure Control Networks. *Review of Applied Science and Technology*, 1(04), 146–182. <https://doi.org/10.63125/5rvjav10>
- [18]. Jahangir, S., & Hammad, S. (2024). A Meta-Analysis of OSHA Safety Training Programs and their Impact on Injury Reduction and Safety Compliance in U.S. Workplaces. *International Journal of Scientific Interdisciplinary Research*, 5(2), 559–592. <https://doi.org/10.63125/8zxw0h59>
- [19]. 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. <https://doi.org/10.1007/s12652-021-03329-z>
- [20]. Jinnat, A., & Md. Kamrul, K. (2021). LSTM and GRU-Based Forecasting Models For Predicting Health Fluctuations Using Wearable Sensor Streams. *American Journal of Interdisciplinary Studies*, 2(02), 32–66. <https://doi.org/10.63125/1p8gbp15>
- [21]. Karniadakis, G. E., Kevrekidis, I. G., Lu, L., Perdikaris, P., Wang, S., & Yang, L. (2021). Physics-informed machine learning. *Nature Reviews Physics*, 3(6), 422–440. <https://doi.org/10.1038/s42254-021-00314-5>
- [22]. King, W. R., & He, J. (2006). A meta-analysis of the technology acceptance model. *Information & Management*, 43(6), 740–755. <https://doi.org/10.1016/j.im.2006.05.003>
- [23]. Kuhar, U., Pantoš, M., Kosec, G., & Švigelj, A. (2018). The impact of model and measurement uncertainties on a state estimation in three-phase distribution networks. *IEEE Transactions on Smart Grid*. <https://doi.org/10.1109/tsg.2018.2823398>
- [24]. Li, X., & Li, F. (2014). GPU-based power flow analysis with Chebyshev preconditioner and conjugate gradient method. *Electric Power Systems Research*, 116, 87–93. <https://doi.org/10.1016/j.epsr.2014.05.005>
- [25]. Lotfifard, S., Kezunovic, M., & Mousavi, M. J. (2011). Voltage sag data utilization for distribution fault location. *IEEE Transactions on Power Delivery*, 26(2), 1239–1246. <https://doi.org/10.1109/tpwrd.2010.2098891>
- [26]. Majidi, M., Etezadi-Amoli, M., Livani, H., & Fadali, M. S. (2016). Distribution systems state estimation using sparsified voltage profile. *Electric Power Systems Research*, 136, 160–168. <https://doi.org/10.1016/j.epsr.2016.02.017>
- [27]. Marangunić, N., & Granić, A. (2015). Technology acceptance model: A literature review from 1986 to 2013. *Expert Systems with Applications*, 42(2), 811–828. <https://doi.org/10.1016/j.eswa.2014.09.020>
- [28]. Mashaly, M. (2021). Connecting the twins: A review on digital twin technology & its networking requirements. *Procedia Computer Science*, 184, 299–305. <https://doi.org/10.1016/j.procs.2021.03.039>
- [29]. Masud, R., & Hammad, S. (2024). Computational Modeling and Simulation Techniques For Managing Rail–Urban Interface Constraints In Metropolitan Transportation Systems. *American Journal of Scholarly Research and Innovation*, 3(02), 141–178. <https://doi.org/10.63125/pxet1d94>
- [30]. Md Ashraful, A., Md Fokhrul, A., & Md Fardaus, A. (2020). Predictive Data-Driven Models Leveraging Healthcare Big Data for Early Intervention And Long-Term Chronic Disease Management To Strengthen U.S. National Health Infrastructure. *American Journal of Interdisciplinary Studies*, 1(04), 26–54. <https://doi.org/10.63125/1z7b5v06>
- [31]. Md Fokhrul, A., Md Ashraful, A., & Md Fardaus, A. (2021). Privacy-Preserving Security Model for Early Cancer Diagnosis, Population-Level Epidemiology, And Secure Integration into U.S. Healthcare Systems. *American Journal of Scholarly Research and Innovation*, 1(02), 01–27. <https://doi.org/10.63125/q8wjee18>
- [32]. Md, K., & Sai Praveen, K. (2024). Hybrid Discrete-Event And Agent-Based Simulation Framework (H-DEABSF) For Dynamic Process Control In Smart Factories. *ASRC Procedia: Global Perspectives in Science and Scholarship*, 4(1), 72–96. <https://doi.org/10.63125/wcqq7x08>
- [33]. Md. Akbar, H., & Farzana, A. (2023). Predicting Suicide Risk Through Machine Learning–Based Analysis of Patient Narratives and Digital Behavioral Markers in Clinical Psychology Settings. *Review of Applied Science and Technology*, 2(04), 158–193. <https://doi.org/10.63125/mqty9n77>
- [34]. Md. Arifur, R., & Haque, B. M. T. (2022). Quantitative Benchmarking of Machine Learning Models for Risk Prediction: A Comparative Study Using AUC/F1 Metrics and Robustness Testing. *Review of Applied Science and Technology*, 1(03), 32–60. <https://doi.org/10.63125/9hd4e011>
- [35]. Md. Towhidul, I., Alifa Majumder, N., & Mst. Shahrin, S. (2022). Predictive Analytics as A Strategic Tool For Financial Forecasting and Risk Governance In U.S. Capital Markets. *International Journal of Scientific Interdisciplinary Research*, 1(01), 238–273. <https://doi.org/10.63125/2rpyze69>
- [36]. Mora-Flórez, J., Meléndez, J., & Carrillo-Caicedo, G. (2008). Comparison of impedance-based fault location methods for power distribution systems. *Electric Power Systems Research*, 78(4), 657–666. <https://doi.org/10.1016/j.epsr.2007.05.010>

- [37]. Mostafa, K. (2023). An Empirical Evaluation of Machine Learning Techniques for Financial Fraud Detection in Transaction-Level Data. *American Journal of Interdisciplinary Studies*, 4(04), 210-249. <https://doi.org/10.63125/60amyk26>
- [38]. Palensky, P. (2022). Digital twins and their use in future power systems. *Digital Twin*, 1, Article 4. <https://doi.org/10.12688/digitaltwin.17435.2>
- [39]. Pan, H., & Liu, C. (2020). *Digital twin and its application in power system* Proceedings of the International Conference on Power and Renewable Energy (ICPRE),
- [40]. Prasad, S., & Kumar, D. M. V. (2018). Robust meter placement for active distribution state estimation using a new multi-objective optimisation model. *IET Science, Measurement & Technology*. <https://doi.org/10.1049/iet-smt.2018.5175>
- [41]. Raghuvamsi, Y., Teeparthi, K., & Kosana, V. (2022). A novel deep learning architecture for distribution system topology identification with missing PMU measurements. *Results in Engineering*, 15, 100543. <https://doi.org/10.1016/j.rineng.2022.100543>
- [42]. 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. <https://doi.org/10.1016/j.jcp.2018.10.045>
- [43]. Rauf, M. A. (2018). A needs assessment approach to english for specific purposes (ESP) based syllabus design in Bangladesh vocational and technical education (BVTE). *International Journal of Educational Best Practices*, 2(2), 18-25.
- [44]. Rifat, C., & Jinnat, A. (2022). Optimization Algorithms for Enhancing High Dimensional Biomedical Data Processing Efficiency. *Review of Applied Science and Technology*, 1(04), 98-145. <https://doi.org/10.63125/2zg6x055>
- [45]. Rifat, C., & Khairul Alam, T. (2022). Assessing The Role of Statistical Modeling Techniques in Fraud Detection Across Procurement And International Trade Systems. *American Journal of Interdisciplinary Studies*, 3(02), 91-125. <https://doi.org/10.63125/gbdq4z84>
- [46]. Rifat, C., & Rebeka, S. (2023). The Role of ERP-Integrated Decision Support Systems in Enhancing Efficiency and Coordination In Healthcare Logistics: A Quantitative Study. *International Journal of Scientific Interdisciplinary Research*, 4(4), 265-285. <https://doi.org/10.63125/c7srk144>
- [47]. Rifat, C., & Rebeka, S. (2024). Integrating Artificial Intelligence and Advanced Computing Models to Reduce Logistics Delays in Pharmaceutical Distribution. *American Journal of Health and Medical Sciences*, 5(03), 01-35. <https://doi.org/10.63125/t1kx4448>
- [48]. Sai Praveen, K. (2024). AI-Enhanced Data Science Approaches For Optimizing User Engagement In U.S. Digital Marketing Campaigns. *Journal of Sustainable Development and Policy*, 3(03), 01-43. <https://doi.org/10.63125/65ebsn47>
- [49]. Salim, R. H., Salim, K. C. O., & Bretas, A. S. (2011). Further improvements on impedance-based fault location for power distribution systems. *IET Generation, Transmission & Distribution*, 5(4), 467-478. <https://doi.org/10.1049/iet-gtd.2010.0446>
- [50]. Salles, D., Coury, D. V., Oleskovicz, M., & Filho, J. M. (2021). Fault localization method for power distribution systems based on gated graph neural networks. *Electrical Engineering*, 103, 2211-2224. <https://doi.org/10.1007/s00202-021-01223-7>
- [51]. Sapountzoglou, N., Lago, J., De Schutter, B., & Raison, B. (2020). A generalizable and sensor-independent deep learning method for fault detection and location in low-voltage distribution grids. *Applied Energy*, 276, 115299. <https://doi.org/10.1016/j.apenergy.2020.115299>
- [52]. Shehwar, D., & Nizamani, S. A. (2024). Power Dynamics in Indian Ocean: US Indo-Pacific Strategic Report and Prospects for Pakistan's National Security. *Government: Research Journal of Political Science*, 13.
- [53]. Shoflul Azam, T., & Md. Al Amin, K. (2024). Quantitative Study on Machine Learning-Based Industrial Engineering Approaches For Reducing System Downtime In U.S. Manufacturing Plants. *International Journal of Scientific Interdisciplinary Research*, 5(2), 526-558. <https://doi.org/10.63125/kr9r1r90>
- [54]. Sifat, M. M. H., Choudhury, S. M., Das, S. K., Ahamed, M. H., Muyeen, S. M., Hasan, M. M., Ali, M. F., Tasneem, Z., Islam, M. M., Islam, M. R., Badal, M. F. R., Abhi, S. H., Sarker, S. K., & Das, P. (2022). Towards electric digital twin grid: Technology and framework review. *Energy and AI*, 11, 100213. <https://doi.org/10.1016/j.egyai.2022.100213>
- [55]. Soltani, Z., & Khorsand, M. (2022). Real-time topology detection and state estimation in distribution systems using micro-PMU and smart meter data. *IEEE Systems Journal*, 16(3), 3554-3565. <https://doi.org/10.1109/jsyst.2022.3153449>
- [56]. Usman, M. U., & Faruque, M. O. (2018). Validation of a PMU-based fault location identification method for smart distribution network with photovoltaics using real-time data. *IET Generation, Transmission & Distribution*, 12(21), 5810-5819. <https://doi.org/10.1049/iet-gtd.2018.6245>
- [57]. Veerakumar, N., Četenović, D., Kongurai, K., Popov, M., Jongepier, A., & Terzija, V. (2022). PMU-based real-time distribution system state estimation considering anomaly detection, discrimination and identification. *International Journal of Electrical Power & Energy Systems*, 148, 108916. <https://doi.org/10.1016/j.ijepes.2022.108916>
- [58]. Venkatesh, V., & Bala, H. (2008). Technology acceptance model 3 and a research agenda on interventions. *Decision Sciences*, 39(2), 273-315. <https://doi.org/10.1111/j.1540-5915.2008.00192.x>
- [59]. Venkatesh, V., Thong, J. Y. L., & Xu, X. (2012). Consumer acceptance and use of information technology: Extending the unified theory of acceptance and use of technology. *MIS Quarterly*, 36(1), 157-178. <https://doi.org/10.2307/41410412>
- [60]. Wang, L., Zhou, Q., & Jin, S. (2020). Physics-guided deep learning for power system state estimation. *Journal of Modern Power Systems and Clean Energy*, 8(4), 607-615. <https://doi.org/10.35833/mpce.2019.000565>

- [61]. Zaman, M. A. U., Sultana, S., Raju, V., & Rauf, M. A. (2021). Factors Impacting the Uptake of Innovative Open and Distance Learning (ODL) Programmes in Teacher Education. *Turkish Online Journal of Qualitative Inquiry*, 12(6).
- [62]. Zamzam, A. S., & Sidiropoulos, N. D. (2020). Physics-aware neural networks for distribution system state estimation. *IEEE Transactions on Power Systems*, 35(6), 4347-4356. <https://doi.org/10.1109/tpwrs.2020.2988352>
- [63]. Zhang, D., Lu, L., Guo, L., & Karniadakis, G. E. (2019). Quantifying total uncertainty in physics-informed neural networks for solving forward and inverse stochastic problems. *Journal of Computational Physics*, 397, 108850. <https://doi.org/10.1016/j.jcp.2019.07.048>
- [64]. Zhang, Z., & Lv, L. (2022). Application status and prospects of digital twin technology in distribution grid. *Energy Reports*, 8, Article in 2022. <https://doi.org/10.1016/j.egyrs.2022.10.410>
- [65]. Zhou, G., Bo, R., Chien, L., Zhang, X., Shi, F., Xu, C., & Feng, Y. (2017). GPU-based batch LU-factorization solver for concurrent analysis of massive power flows. *IEEE Transactions on Power Systems*, 33(1), 1123-1131. <https://doi.org/10.1109/tpwrs.2017.2662322>
- [66]. Zhou, G., Bo, R., Chien, L., Zhang, X., Yang, S., & Su, D. (2018). GPU-accelerated algorithm for online probabilistic power flow. *IEEE Transactions on Power Systems*, 33(1), 1132-1135. <https://doi.org/10.1109/tpwrs.2017.2756339>
- [67]. Zhou, G., Feng, Y., Bo, R., Chien, L., Zhang, X., Lang, Y., Jia, Y., & Chen, Z. (2017). GPU-accelerated batch-ACPF solution for N-1 static security analysis. *IEEE Transactions on Smart Grid*, 8(6), 2841-2850. <https://doi.org/10.1109/tsg.2016.2600587>
- [68]. Zhou, G., Hua, J., Zhao, J., Feng, Y., Yao, Y., & Fu, M. (2021). GPU-based matrix structure driven state estimation for large-scale power systems. *International Journal of Electrical Power & Energy Systems*, 133, 107298. <https://doi.org/10.1016/j.ijepes.2021.107298>