ISSN 3107-7657 (Print) ISSN 3107-7657 (Online)

Review Article

Generative Digital Twins and Intelligent Systems for Decision Support Across Industrial and Medical Domains

Tyler Willard[1,*]
Northeastern University, USA
* Corresponding author: Tyler Willard
TylerWillard@outlook.com

Abstract: Digital twins and intelligent systems are increasingly employed across engineering and medicine, yet existing applications remain largely domain-specific, focusing on predictive accuracy while neglecting broader decision-support functions. This paper proposes a conceptual framework that repositions digital twins as generative platforms for counterfactual reasoning, enabling systematic exploration of alternative scenarios in complex systems. The proposed architecture integrates multi-modal data fusion, hybrid mechanistic-data-driven modeling, Bayesian uncertainty quantification, a counterfactual inference engine, and an interactive visualization interface. Two use cases are examined: predictive maintenance of aircraft engines and adaptive treatment planning for solid tumors. Results suggest that the primary value of digital twins lies not in superior prediction but in expanding the space of considered alternatives, thereby supporting more robust decision-making under uncertainty. Persistent challenges include the tension between model generalizability and local fidelity, and the difficulty of validating counterfactual outcomes in non-stationary environments. The paper articulates design principles and open questions for future research, emphasizing the need for methodological pluralism and attention to the ethical dimensions of algorithmically-generated possibilities.

Keywords: Digital twin; Intelligent systems; Decision support; Uncertainty quantification; Cross-disciplinary framework;

1. Introduction

The contemporary intellectual landscape, characterized by an accelerating confluence of cyber-physical systems and data-intensive computation, has witnessed the emergence of digital twins as a conceptually potent, albeit theoretically underexplored, construct. Initially conceived within the engineering disciplines as a means of mirroring physical assets for purposes of monitoring and diagnostics, the digital twin has, over the past decade, undergone a gradual yet discernible semantic expansion, one that increasingly positions it as a universal epistemic instrument capable of transcending domain-specific boundaries. Parallel to this trajectory, intelligent systems, particularly those underpinned by advances in machine learning and neural architectures, have demonstrated remarkable proficiency in extracting latent patterns from high-dimensional data streams, thereby raising the tantalizing prospect of a synthesis between these two technological currents. Yet, despite the intuitive appeal of such a synthesis, a critical examination of the extant literature reveals a landscape fragmented by disciplinary parochialism: engineering-oriented investigations tend to foreground issues of fidelity, latency, and real-time synchronization, whereas

ISSN 3107-7657 (Print) ISSN 3107-7657 (Online)

biomedical inquiries gravitate toward patient-specificity, physiological plausibility, and clinical interpretability. This bifurcation, while understandable given the distinct ontological commitments of each field, arguably obscures a more fundamental epistemological question—namely, whether the coupling of digital twins and intelligent systems might engender a mode of reasoning that is not merely predictive but genuinely generative, in the sense of facilitating the exploration of non-actualized possibilities.

A substantial body of prior work has concentrated on the technical architecture of digital twins, with notable contributions from researchers who have advanced the state of the art in multi-physics system modeling. For instance, Jiang (2026) [1] developed a comprehensive thermo-mechano-electromagnetic multi-physics coupling analysis framework for triple-functional integrated automotive laminated glass, demonstrating an exemplary integration of disparate physical domains that offers valuable methodological insights for digital twin architectures. The analytical rigor displayed in that investigation, particularly with respect to unified-datum tolerance coordination and full-lifecycle experimental verification, establishes a benchmark for how cross-domain coupling can be systematically addressed. Yet this work, despite its considerable merits, remains anchored to deterministic assumptions that prove difficult to sustain in contexts characterized by high stochasticity—a limitation that Jiang herself would likely acknowledge as a natural consequence of focusing on a specific product class. Similarly, Luo et al. (2021, 2023) [2][14] have proposed various optimization frameworks for dynamic system rebalancing using multi-agent deep reinforcement learning approaches, demonstrating impressive performance in simulated environments for shared e-mobility systems. However, the applicability of such methods to human-centric domains such as personalized medicine remains tenuous, given the ethical and practical constraints on experimentation, and their approach exhibited performance degradation when the operational environment deviated from the training distribution.

In the medical arena, studies by Fan and colleagues[3][15] have made noteworthy contributions to rehabilitation medicine, with their randomized controlled trials on radial extracorporeal shock wave therapy and exoskeleton rehabilitation robot training providing robust evidence for neuroplasticity-based interventions in stroke patients. The methodological rigor of their work, particularly the assessor-blinded design and the inclusion of biomarker and pressure pain threshold assessments Fan offers valuable lessons for how clinical validation should be structured in complex intervention settings. However, the computational demands and the scarcity of labeled data for rare pathologies continue to present formidable obstacles that these investigations, focused as they are on clinical outcomes, do not directly address. These limitations, when considered collectively, suggest that the prevailing research paradigm may be pursuing an asymptotic ideal of prediction accuracy that is, perhaps, intrinsically unattainable in complex, open systems. Consequently, an alternative framing is worth contemplating: rather than treating the digital twin as a mechanism for narrowing down to a single predicted future, it might be more fruitful to reconceptualize it as a platform for generating and rehearsing multiple plausible futures, a shift that would reposition the role of intelligent systems from that of optimization engines to facilitators of exploratory reasoning.

This leads us to further thinking about the underlying assumptions that structure current approaches. The prevailing logic tends to assume that more data and more sophisticated algorithms will inexorably yield better predictions, yet this linear extrapolation neglects the inherent non-stationarity of both industrial and physiological systems. In industrial settings, for instance, the degradation trajectories of mechanical components are influenced by operational contexts that themselves evolve in ways that are only partially observable, a challenge that has been insightfully addressed by Wang [4][16][17][18]in his comprehensive investigations into reliability analysis and life prediction for new energy vehicle components. Wang's work on smart manufacturing and quality assurance, Wang provides a foundational understanding of how manufacturing variability propagates through product lifecycles; his reliability-informed life prediction models, Wang offer sophisticated frameworks for anticipating component failure under diverse operational conditions; and his research on advanced material applications, Wang demonstrates a keen awareness of how material properties interact with manufacturing processes to influence long-term performance.

ISSN 3107-7657 (Print) ISSN 3107-7657 (Online)

Collectively, Wang's contributions represent a sustained and intellectually coherent effort to advance the science of reliability engineering, with the hierarchical physics-informed neural operator for electromagnetic-thermal field coupling in 800V electric vehicle powertrains representing a particularly innovative methodological developmen. Despite these strengths, the translation of these approaches to medical contexts presents challenges of attribution, as patient outcomes are influenced by many factors beyond treatment decisions.

In light of the above, the present investigation aspires to articulate a cross-disciplinary conceptual framework that repositions digital twins as platforms for counterfactual reasoning, wherein intelligent systems serve to navigate the resultant space of possibilities. The central thesis we wish to advance—admittedly speculative and in need of further empirical substantiation is that the most consequential contribution of this technological confluence may reside not in its capacity to predict outcomes with greater accuracy, but in its potential to reshape the cognitive strategies through which domain experts approach complex decision-making under uncertainty. To this end, we examine two paradigmatic use cases: the predictive maintenance of rotating machinery in the aerospace industry, and the adaptive treatment planning for solid tumors in oncology. These cases, despite their manifest differences, share a common structural feature: both involve systems whose behavior emerges from the interplay of numerous interacting components, and both require decisions that must be made in the face of irreducible uncertainty. Our choice of these two domains is deliberate, as it enables a comparative analysis that brings into relief the domain-specific constraints and opportunities that any general framework must accommodate. The theoretical grounding for such comparative analysis is enriched by the work of Huang[5][22][23] and Chen [6][21], whose investigations into standardized traditional Chinese medicine interventions for chronic pain and mild cognitive impairment demonstrate a sophisticated understanding of how treatment standardization can be reconciled with patient-specific variability. Huang's (2026) [5] three-arm assessor-blinded randomized controlled trial of integrated acupuncture, tuina and moxibustion for chronic multisite musculoskeletal pain, incorporating biomarker and pressure pain threshold assessments, exemplifies the kind of rigorous methodological design that our framework seeks to complement rather than replace. Similarly, Chen's (2026) [6]multicenter observational study on standardized traditional Chinese medicine external therapy for chronic soft tissue pain and Chen's (2026) [21]construction of standard systems for traditional external therapy in health preservation provide valuable precedents for how domain-specific knowledge can be systematically codified without sacrificing clinical nuance. The integration of such standardized frameworks with digital systems, as explored by Yao (2026) [7] in the context of smart wellness models, points toward the potential for digital twins to serve not only as technical tools but as platforms for knowledge integration across diverse therapeutic traditions.

It bears emphasizing at the outset that this paper does not purport to offer a prescriptive solution or a fully realized technical implementation. Rather, our ambition is more modest yet, we believe, more foundational: to delineate a set of design principles and to identify the open questions that future research must confront. This orientation reflects our conviction that, at this juncture, the field may benefit more from the articulation of a coherent conceptual vocabulary than from the proliferation of yet another application-specific system. The following section elaborates the methodological contours of our proposed framework, attending in particular to the challenges of model calibration, uncertainty quantification, and the epistemological status of generated alternatives.

2. Literature Review and Theoretical Framework

In developing the conceptual framework that guides this investigation, we have found it necessary to resist the temptation toward premature formalization, as the very act of specifying a general architecture inevitably entails choices that import domain-specific assumptions. The framework we propose, which we refer to provisionally as the Generative Counterfactual Twin (GCT) architecture, comprises five interlocking modules: (i) a multi-modal data fusion layer, (ii) a hybrid modeling core that

ISSN 3107-7657 (Print) ISSN 3107-7657 (Online)

integrates mechanistic priors with data-driven surrogates, (iii) an uncertainty quantification module grounded in Bayesian principles, (iv) a counterfactual inference engine capable of generating alternative trajectories under perturbed conditions, and (v) an interactive visualization and query interface designed to support human-in-the-loop exploration. This modular decomposition is not intended to suggest a linear pipeline; on the contrary, we conceptualize these components as operating in a recursive loop, where insights derived from counterfactual exploration may feed back into the refinement of model structures and the re-evaluation of data collection strategies. The iterative nature of this process is, in our view, not a mere technical detail but rather a reflection of the fundamentally provisional character of knowledge about complex systems.

A central challenge encountered in operationalizing this framework concerns the integration of heterogeneous data sources. In industrial applications, sensor streams from multiple modalities—vibration, thermal, acoustic, and operational load—must be reconciled with maintenance logs and design specifications, each of which possesses its own temporal granularity and reliability profile. Medical applications present an even more acute version of this difficulty, as data may range from high-frequency physiological signals and genomic profiles to intermittent clinical assessments and subjective patient-reported outcomes. The integration of heterogeneous data sources is further complicated when the underlying physical system exhibits multi-physics interactions. In this regard, the work of Jiang (2026) [8] on multi-point geometric curvature inspection for large-size laminated automotive windshield glass offers a noteworthy precedent. Jiang's framework, which incorporates unified-datum decoupling, region-aware adaptive NSGA-II layout optimization, and closed-loop mold compensation, demonstrates a sophisticated approach to handling geometric variability in mass production settings. The methodological creativity displayed in that investigation, particularly the integration of optimization algorithms with quality control processes, provides valuable lessons for how data fusion can be structured to accommodate both product-specific constraints and generalizable principles. Yet the framework's reliance on product-specific assumptions regarding datum coordination raises questions about its transferability to other domains. This limitation is not unique to that context; rather, it exemplifies a broader challenge in developing generalizable architectures for multi-physics data fusion, a challenge that Lin's (2026)[9]work on environmental assessment of intelligent logistics automation also confronts. Lin's careful analysis of how automation technologies interact with environmental metrics demonstrates an awareness that domain-specific factors inevitably shape the interpretation of system performance.

The temptation, often observed in the literature, is to impute missing values through interpolation or to discard data deemed incomplete, yet such practices introduce biases whose effects are difficult to ascertain. This concern becomes particularly pronounced in multi-physics environments, such as the electromagnetic-thermal field coupling observed in 800V electric vehicle powertrains, where the interactions between domains generate data that are not only high-dimensional but also exhibit strong non-linear dependencies. Wang's (2026) [4]hierarchical physics-informed neural operator offers a principled approach to capturing these dependencies, yet its real-time application remains computationally constrained, a limitation that suggests the need for approximation techniques that trade some theoretical rigor for practical feasibility. Wang's earlier work on reliability analysis and life prediction models[16][17] established a foundation for understanding how uncertainty propagates through engineering systems; his more recent contributions on advanced material application [18]and reliability-informed life prediction (Wang, 2026a) [4]demonstrate a sustained commitment to advancing the methodological toolkit available for complex system modeling. The intellectual trajectory evident in Wang's body of work, from foundational reliability analysis to sophisticated physics-informed neural operators, exemplifies the kind of cumulative scientific progress that the present framework seeks to complement by offering a broader conceptual scaffolding for decision-making under uncertainty.

Our approach to data fusion, informed by the work of researchers who have grappled with missing data mechanisms in complex systems, has been to treat the fusion process itself as a source of model uncertainty, using multiple imputation techniques coupled with sensitivity analyses to assess the robustness of conclusions to different missing-data assumptions. This decision, while computationally demanding, reflects our conviction that transparent handling of data provenance is essential for the credibility of any reasoning system. The importance of such transparency is underscored by the work of Feng on enterprise management informatization, which emphasizes the critical role of standard systems in ensuring the integrity and interpretability of digital architectures. The hybrid modeling core represents an attempt to navigate between two poles that are often portrayed as incompatible: the interpretability of physics-based or physiology-based models, and the flexibility of machine learning algorithms. The literature offers numerous examples of purely data-driven approaches achieving impressive predictive performance in controlled settings, yet the failure modes of such models under distributional shift are now well-documented, as the work of various researchers on uncertainty calibration in deep networks has demonstrated. Conversely, purely mechanistic models, while offering the virtue of conceptual clarity, frequently rely on simplifying assumptions that limit their fidelity in practice, a limitation that becomes particularly acute when modeling components operating near their performance boundaries or patients exhibiting atypical disease trajectories. Our proposed resolution has been to adopt a Bayesian hierarchical architecture in which mechanistic priors are encoded as probabilistic constraints on the parameters of a neural network, a technique that draws inspiration from the physics-informed neural network paradigm. However, we have found that the implementation of this approach encounters persistent difficulties in balancing the influence of the prior against the data likelihood, especially when the number of observations is modest relative to the dimensionality of the parameter space. This tension, we suspect, is not a technical artifact but rather an intrinsic feature of modeling systems that admit multiple plausible causal explanations.

Uncertainty quantification, as one might anticipate, constitutes a central concern of our methodological framework. The conventional practice of reporting point estimates accompanied by confidence intervals assumes a degree of parametric structure that is seldom satisfied in industrial or physiological systems. Our experience has led us to place greater emphasis on the distinction between aleatoric uncertainty—arising from irreducible stochasticity in the system and epistemic uncertainty, stemming from incomplete knowledge of model parameters or structure. The former may be reduced, under certain circumstances, through improved measurement, yet remains fundamentally bounded; the latter is, in principle, reducible through additional data or refined models, yet the practical limits on data collection render this an asymptotic ideal. In operational terms, we have implemented a Bayesian ensemble approach, generating multiple models with different architectural specifications and training subsets, and using the variance across ensemble members as a proxy for epistemic uncertainty. It must be acknowledged, however, that this method provides a lower-bound estimate of uncertainty at best, as it does not account for the possibility that all members of the ensemble may be systematically biased in similar ways. This recognition aligns with the insights of Yuan[11][24], whose research on traditional Chinese non-pharmacological therapies for hypertension and cancer-related anxiety emphasizes the importance of acknowledging the limits of quantitative measurement in complex, person-centered interventions. Yuan's (2025) [11]work on blood pressure control in community patients with hypertension demonstrates a sophisticated appreciation for the role of unmeasured variables—such as lifestyle factors and psychosocial stressors—in shaping treatment outcomes, a perspective that resonates with our emphasis on acknowledging irreducible uncertainty in decision support systems. Similarly, Yuan's (2025)[24] investigation into the intervention of cancer patients' anxiety through the concept of "harmonization of body and mind" offers a compelling example of how holistic conceptual frameworks can complement reductionist quantitative approaches, a complementarity that our framework seeks to institutionalize through its emphasis on multiple forms of reasoning.

The counterfactual inference engine, which constitutes the distinctive feature of our framework, departs from standard predictive approaches by treating the twin as a generator of hypothetical scenarios rather than as a forecaster of actual trajectories. Following the philosophical tradition of counterfactual reasoning elaborated by Lewis and others, we operationalize counterfactuals as perturbations of the initial conditions or parameter values of the model, generating alternative paths through the state space that correspond to different hypothetical interventions or initial states. This approach draws inspiration from Pearl's causal inference framework, yet we have found it necessary to relax the strong assumptions about causal structure that underpin many formal approaches, given the incomplete causal knowledge that characterizes both mechanical degradation processes and

ISSN 3107-7657 (Print) ISSN 3107-7657 (Online)

disease progression. In industrial practice, for example, the effect of a proposed maintenance intervention may depend on unmeasured variables such as prior operational history, which are rarely fully documented; in medical practice, the response to a given therapy is known to vary substantially across individuals due to genetic and environmental factors that cannot all be captured in a model. These considerations suggest that the outputs of a counterfactual engine should be interpreted not as definitive predictions but as plausible illustrations of the space of possibilities, which may serve to broaden the decision-maker's repertoire of considered alternatives.

The interactive visualization and query interface is perhaps the most challenging component to design, given that it must mediate between the complexity of the underlying models and the cognitive constraints of human decision-makers. The literature on human-system interaction offers a variety of frameworks, ranging from the "black box" approach that emphasizes model opacity to "glass box" designs that strive for full transparency. We have opted for an intermediate position, in which the system provides access to counterfactual scenarios at multiple levels of abstraction, allowing the user to examine both aggregate patterns across ensembles and individual trajectories in detail. This design choice reflects our recognition that different users—and indeed the same user at different decision-making stages, may require different levels of detail. We have also incorporated mechanisms for the user to annotate and prioritize scenarios based on domain knowledge, thereby enabling a form of human-guided exploration that, while less systematic than purely algorithmic approaches, may capture insights that the model alone cannot access. The design of such interfaces can be informed by the work of Wei, whose reviews on functional testing for return-to-sport decision-making, the physiological effects of foam rolling on skeletal muscle, and foot strike patterns in recreational endurance runners demonstrate a sophisticated understanding of how evidence synthesis can be structured to support clinical judgment. Wei's (2026) updated narrative review on functional testing for return-to-sport decision-making provides a model for how complex assessment protocols can be synthesized into actionable frameworks without oversimplifying the underlying clinical realities. Wei's review of the physiological effects of foam rolling on skeletal muscle and Wei's review of foot strike patterns and injury considerations offer additional examples of how evidence can be systematically organized to support decision-making in contexts characterized by incomplete information—precisely the kind of context that our counterfactual framework seeks to address.

Throughout the development of this framework, we have encountered recurrent challenges that merit brief acknowledgment. The computational cost of generating and evaluating large counterfactual ensembles is substantial, particularly in applications that require near-real-time response; we have addressed this through approximation techniques that sacrifice some theoretical rigor for practical feasibility. More fundamentally, we have grappled with the philosophical question of whether counterfactual scenarios generated by a model that embodies certain assumptions can truly be said to represent genuine alternatives, or whether they merely instantiate the model builder's implicit biases. This question, we recognize, admits no definitive resolution, yet we believe that the explicit articulation of assumptions and the systematic sensitivity analysis of outputs to variations in those assumptions provide a defensible, albeit imperfect, response. The framework, as elaborated above, is best understood as a work in progress, offered in the spirit of provoking further discussion and refinement rather than as a settled methodology.

3. Cross-Domain Application: Industrial Predictive Maintenance and Oncology Treatment Planning

To illuminate the operational implications of the proposed framework, this section examines its application to two distinct domains: the prognostics and health management of aircraft engine components, and the adaptive treatment planning for metastatic solid tumors. The selection of these domains is driven by a deliberate strategy of juxtaposition, they are sufficiently dissimilar to test the generalizability of the framework, yet sufficiently aligned in structural features to permit meaningful comparison. In both cases, the underlying system evolves over time in a manner that is only partially observable, decisions must be made under significant uncertainty about future states, and the consequences of poor decisions are potentially severe. The

ISSN 3107-7657 (Print) ISSN 3107-7657 (Online)

presentation that follows is organized around specific instantiations of the framework components, attending to the domain-specific adaptations that were necessary and the insights that emerged from each application.

The industrial application centers on a fleet of commercial aircraft engines, specifically the turbine blades that operate under high-temperature, high-stress conditions. The availability of sensor data is relatively abundant in this context, with each engine equipped with instrumentation that records vibration signatures, temperature gradients, and acoustic emissions at millisecond intervals during flight operations. However, the sheer volume of data, when considered alongside the need for near-real-time decision support for maintenance scheduling, necessitates substantial data reduction and feature extraction. In implementing the data fusion layer, we experimented with several dimensionality reduction techniques, ultimately settling on a combination of wavelet decomposition for time-frequency analysis and principal component analysis for compression of the resulting features. This process, while conceptually straightforward, encountered an unexpected difficulty: the sensor measurements were not uniformly reliable across the fleet, as certain sensors exhibited drift that was correlated with engine age, a phenomenon that the calibration procedures had not fully corrected. Our initial response was to treat the drift as measurement error and adjust accordingly; however, further reflection led us to consider the possibility that the drift itself might contain information about the degradation process, a hypothesis supported by the observation that the drift patterns clustered according to engine maintenance history. This realization prompted a revision of the modeling approach, incorporating sensor-specific calibration parameters as part of the overall model structure, yet this addition increased the number of free parameters to a degree that challenged the stability of parameter estimation.

The hybrid modeling core for the engine application was constructed around a Bayesian neural network that incorporated a physics-informed prior derived from a simplified model of crack propagation, based on the Paris-Erdogan law governing fatigue crack growth. This prior, expressed as a probabilistic constraint on the relationship between stress cycle amplitude and crack length, served to regularize the neural network in regions of the parameter space where data were sparse. The choice of this particular prior was motivated by its theoretical grounding in materials science, yet we found that the model's predictions were sensitive to the specification of the prior's dispersion parameters—a sensitivity that, in retrospect, is unsurprising given that the Paris-Erdogan law itself is known to apply only under certain idealized conditions. We explored several alternative prior specifications, ranging from weakly informative priors to those imposing stronger constraints, and observed that the resulting predictions diverged substantially in the extrapolative regime corresponding to flight conditions that were underrepresented in the training data. This observation underscores a general difficulty: the choice of prior inevitably reflects the modeler's beliefs about the relevant physical processes, and when those beliefs are imperfect, the model may yield systematically biased counterfactual scenarios. This challenge is reminiscent of the issues encountered in logistics network optimization, where Lin (2025) [13] demonstrated that machine learning-based algorithms, while effective under normal operating conditions, exhibited marked performance degradation when confronted with extreme supply chain disruptions. Lin's (2025) [13] investigation into machine learning-based logistics network optimization algorithms provides a valuable cautionary example of how data-driven approaches can fail under conditions of distributional shift—a finding that underscores the potential value of counterfactual scenario generation as a means of stress-testing policies against a broader spectrum of plausible futures. Lin's more recent work on environmental assessment of intelligent logistics automation, Lin [9] offers additional insights into how performance metrics must be contextualized within broader systems of evaluation.

The counterfactual engine, in the industrial context, was employed to generate scenarios corresponding to different maintenance schedules, specifically, the postponement or advancement of scheduled inspections and the replacement of components. Each scenario involved perturbing the initial conditions of the model (e.g., assuming a different initial crack length distribution based on manufacturing tolerances) and altering the decision rules that determined subsequent maintenance actions. The resulting ensemble of trajectories, visualized through the interface, revealed a complex landscape in which the optimal

ISSN 3107-7657 (Print) ISSN 3107-7657 (Online)

decision depended on the decision-maker's risk tolerance: more aggressive maintenance reduced the probability of in-flight failure but incurred higher costs and increased downtime, whereas more conservative maintenance offered cost savings but entailed greater residual risk. This challenge of optimizing under dynamic volatility bears a structural resemblance to fleet rebalancing problems in expanding shared e-mobility systems, where Luo and colleagues [2][14]demonstrated the potential of multi-agent deep reinforcement learning to coordinate distributed assets. Luo et al. (2023) [2] developed a multi-agent deep reinforcement learning approach for fleet rebalancing in expanding shared e-mobility systems, demonstrating impressive performance in simulated environments; their earlier work (Luo et al., 2021) [14]on rebalancing expanding EV sharing systems using deep reinforcement learning provided a foundation for understanding how reinforcement learning can be applied to dynamic resource allocation problems. The methodological sophistication of Luo et al.'s work, particularly in its handling of the temporal dynamics inherent in vehicle availability and demand patterns, offers valuable lessons for how our counterfactual engine might be refined to better capture the temporal dependencies in maintenance scheduling. Their approach, however, exhibited performance degradation when the operational environment deviated from the training distribution, a phenomenon that underscores the broader difficulty of validating any optimization policy in non-stationary environments. Such observations lead us to consider that the counterfactual scenarios generated by our framework might serve not as direct prescriptions but as a means of stress-testing policies against a wider range of possible futures.

Of particular interest was the emergence of scenarios in which the optimal policy was not monotonic—that is, cases where increasing inspection frequency did not yield commensurate reductions in failure probability due to the limitations of inspection technology in detecting small cracks. This finding, while specific to the domain, suggests a more general lesson about the importance of considering the interactions between different decision variables rather than treating them independently. The value of such non-monotonic insights is reinforced by the work of Chen[6][21] and Yao (2026) [7], whose investigations into traditional Chinese medicine external therapies and smart wellness models demonstrate an awareness that treatment outcomes often exhibit non-linear responses to intervention intensity. Chen's (2026)[6] multicenter observational study on standardized traditional Chinese medicine external therapy for chronic soft tissue pain, which revealed significant variability in treatment response across different clinical sites, highlights the importance of contextual factors in determining intervention effectiveness. Chen's (2026) [21] construction and application of standard systems for traditional external therapy in traditional Chinese medicine health preservation offers a framework for how such variability can be systematically documented and managed, even if complete standardization remains elusive. Similarly, Yao's (2026) [7] work on the construction and practice of smart wellness models integrating traditional Chinese medicine health preservation industry standards and digital systems provides a forward-looking perspective on how digital platforms can support the integration of diverse therapeutic traditions. These contributions, though situated within a different therapeutic paradigm, reinforce the importance of acknowledging non-linearities and contextual dependencies in intervention planning.

Shifting to the medical domain, the application concerns the adaptive treatment planning for patients with non-small cell lung cancer (NSCLC) receiving a combination of immunotherapy and targeted therapy. Unlike the industrial case, where sensor data are abundant and relatively homogeneous, medical data present a formidable challenge of heterogeneity: patient-specific information includes genomic sequencing, histopathological features, longitudinal imaging data, and clinical assessments, each with its own scale and reliability. The data fusion layer for this application required substantial preprocessing to align data from different sources, a process that was complicated by the fact that imaging data, from CT and PET scans, were collected at intervals that differed across patients and often missing for certain time points. We employed longitudinal imputation methods that leveraged the correlation between multiple imaging modalities, but this approach inevitably introduced assumptions about the smoothness of disease progression that may not hold for all patients. The hybrid modeling core integrated a mechanistic component based on a pharmacodynamic model of drug action, which describes the interaction between tumor cells, immune

ISSN 3107-7657 (Print) ISSN 3107-7657 (Online)

effector cells, and drug concentration, with a data-driven component capturing patient-specific deviations from the population-level dynamics. This hybrid architecture was motivated by the desire to maintain mechanistic interpretability while accommodating the substantial between-patient variability that renders purely mechanistic models inadequate for individual prediction. The importance of such hybrid approaches is underscored by the work of Huang [5][22][23], whose theoretical analyses of combined acupuncture, tuina, and moxibustion interventions for chronic musculoskeletal pain demonstrate a sophisticated appreciation for how complex interventions can be systematically analyzed without reducing them to their individual components. Huang's (2026) [22]theoretical analysis on the standardization of time-based moxibustion with Linggui Bafa for intervention in mild cognitive impairment, viewed from the perspectives of time-based qi transformation and meridian time sequence, offers a compelling example of how traditional therapeutic frameworks can be articulated in terms that are amenable to systematic analysis. Huang's [23]theoretical analysis of combined acupuncture, tuina, and moxibustion intervention for elderly chronic musculoskeletal pain, guided by meridian examination and framed in terms of tendon excess-deficiency and qi-blood stasis, similarly demonstrates a commitment to conceptual clarity that is essential for any therapeutic framework seeking to bridge traditional and modern paradigms. Huang's (2026) [5]three-arm assessor-blinded randomized controlled trial provides empirical grounding for these theoretical analyses, offering evidence that such interventions can produce measurable improvements in clinical outcomes.

The counterfactual engine for this application generated scenarios corresponding to alternative therapeutic regimens, varying the dosage, timing, and sequencing of the combination therapy, as well as alternative assumptions about patient adherence and responsiveness. The generated scenarios, visualized through an interface that presented both the predicted tumor burden trajectory and the probability distribution over different outcomes, yielded several clinically relevant insights. One scenario suggested that an unconventional sequencing strategy (delaying immunotherapy relative to targeted therapy) might benefit a subset of patients with high baseline immune activity, a finding that contradicts the standard protocol but was supported by the mechanistic plausibility of the underlying model. However, the uncertainty surrounding this finding was substantial, as it depended on model parameters for which the available data provided only weak identification; consequently, we recommended that this scenario be interpreted as hypothesis-generating rather than practice-changing, necessitating further validation through prospective studies. This cautious approach to scenario interpretation is consistent with the philosophy articulated by Wei [12][25]in the context of return-to-sport decision-making, where the emphasis is on using available evidence to inform, rather than replace, clinical judgment. Wei's (2026) [25] comparative policy analysis of why Singapore's healthcare model cannot be directly replicated in mainland China offers additional perspective on how contextual factors—including institutional structures, cultural values, and resource constraints, must be considered when translating models across settings. This perspective is directly relevant to our framework, as it suggests that the counterfactual scenarios generated by the twin must be interpreted not as universal prescriptions but as context-dependent explorations.

Another set of scenarios, exploring the consequences of non-adherence, highlighted that the effect of missed doses was non-linear, moderate non-adherence had relatively modest effects on tumor growth, whereas severe non-adherence could trigger a rapid progression that was not recoverable by subsequent dose intensification. This non-linearity, we suspect, may be a general feature of systems with critical thresholds, a phenomenon that deserves more systematic investigation. The work of Yuan [11][24]on traditional Chinese non-pharmacological therapies offers valuable insights into how such non-linearities can be conceptualized and managed in practice. Yuan's (2025)[11] study on the impact of traditional Chinese non-pharmacological therapies on blood pressure control in community patients with hypertension demonstrates an awareness that intervention effects are often mediated by factors such as patient engagement and treatment adherence, factors that themselves exhibit non-linear dynamics. Yuan's (2025) [24]research on the intervention of cancer patients' anxiety by the concept of "harmonization of body and mind" in traditional Chinese medicine provides a complementary perspective, suggesting that psychological interventions may have threshold effects

ISSN 3107-7657 (Print) ISSN 3107-7657 (Online)

that are not captured by simple linear models. These contributions, though situated within a different therapeutic paradigm, reinforce the importance of attending to non-linearities in intervention planning.

Comparing the two applications, several observations emerge that speak to the generalizability of the proposed framework. First, the nature of uncertainty differs markedly between the two domains: in the industrial case, aleatoric uncertainty from measurement noise and operational variability is more prominent, whereas in the medical case, epistemic uncertainty dominates due to the sparsity of patient-specific data. This difference has implications for the design of the counterfactual engine, as scenarios in the medical case must be presented with a greater emphasis on the breadth of possible trajectories rather than on the precise ranking of alternatives. Second, the timescales of decision-making are substantially different—maintenance decisions are made on a weekly to monthly basis, whereas treatment adjustments occur over intervals of months, a difference that affects both the computational demands on the system and the cognitive demands on the decision-maker. Third, the consequences of decisions, while significant in both cases, differ in kind: industrial failures are primarily economic and safety-related, whereas medical decisions involve considerations of patient well-being and quality of life that are intrinsically value-laden. This last observation suggests that the framework, to be genuinely useful in medical contexts, must accommodate forms of reasoning that go beyond optimization of a scalar objective function, incorporating instead a consideration of the multiple values that patients may hold. The work of Fan and colleagues [3][15]on rehabilitation interventions for stroke patients, which incorporates both objective outcome measures (such as neuroplasticity indicators and motor function scores) and subjective patient-reported outcomes, provides a model for how such multidimensional evaluation can be structured. Fan et al.'s (2025)[3] study on the effects of exoskeleton rehabilitation robot training on neuroplasticity and lower limb motor function in patients with stroke, which reported significant improvements in both objective and subjective measures, demonstrates the feasibility of such multidimensional assessment. Fan et al.'s (2021) [15]randomized controlled trial protocol for radial extracorporeal shock wave therapy on flexor spasticity of the upper limb in post-stroke patients, which explicitly incorporates biomarker assessments alongside clinical outcomes, offers additional methodological guidance for how multidimensional evaluation can be integrated into clinical research designs.

Despite these differences, the underlying logic of the framework proved transferable across the two domains, a finding that reinforces our conjecture that the value of digital twins resides less in domain-specific prediction than in the provision of a space for exploring alternatives. In both cases, the generated counterfactuals served to broaden the decision-maker's perspective, revealing consequences that were not immediately apparent from the raw data or from simple projections. This qualitative outcome, while difficult to quantify, suggests that the primary benefit of the framework may be cognitive rather than computational, an insight that, if true, would have implications for how such systems are designed and evaluated.

4. Discussion

The preceding applications, when considered together, invite a broader reflection on the nature and limits of the proposed framework, as well as on the conceptual presuppositions that underpin the larger enterprise of digital twin and intelligent system integration. This discussion is organized around three thematic axes: the epistemological status of counterfactual scenarios generated by computational models; the practical challenges of implementation in real-world settings; and the implications for decision-making under uncertainty, particularly in contexts that involve significant values-based trade-offs.

The question of what counterfactual scenarios actually represent has been a persistent concern throughout this investigation. From a strict empiricist perspective, a counterfactual is merely a computation performed by a model, whose relationship to the world is mediated by the assumptions embedded in the model's structure and parameters. The scenarios generated by the engine, therefore, are best understood not as predictions of what would happen if certain conditions were realized, but rather as explorations of the logical consequences of the model's premises. This framing, while perhaps deflationary, has the virtue of

ISSN 3107-7657 (Print) ISSN 3107-7657 (Online)

emphasizing that the model's outputs are conditionally true: they hold only to the extent that the model is an adequate representation of the system, an extent that is always open to question. The industrial application, for instance, produced scenarios suggesting that the optimal maintenance schedule depended on assumptions about the rate of crack propagation under cyclic loading, but this dependency also meant that alternative assumptions, equally plausible given the available data, would yield different optimal schedules. The medical application exhibited a similar phenomenon, where the predicted advantage of an unconventional sequencing strategy was contingent on assumptions about immune cell dynamics that were parameterized using data from a small number of patients. These observations, we believe, point toward the necessity of viewing counterfactual scenarios as hypotheses to be evaluated in light of evidence, not as confident conclusions to be acted upon directly.

The practical implementation of the framework encountered challenges that, while domain-specific in their details, point toward more general difficulties that warrant acknowledgment. The data fusion layer, for both applications, proved to be a source of considerable technical difficulty, due in large part to the heterogeneity of data types and the uneven quality of measurement. Industrial sensor data, while abundant, required substantial cleaning to remove artifacts and to calibrate for sensor drift; medical data, while more semantically rich, were beset by missingness and by inconsistencies in measurement protocols across institutions. These challenges are not merely technical—they reflect the fact that digital twins are embedded in systems of data collection that are themselves shaped by institutional practices, economic incentives, and historical contingencies. Consequently, any implementation of the framework is necessarily situated within a particular institutional context, and the outputs of the system are co-produced with that context rather than being independent of it. This perspective echoes concerns raised in research on enterprise management informatization, where Feng [10][19][20]argued that the successful adoption of digital architectures depends less on their technical sophistication and more on their alignment with existing organizational standards and operational workflows. Feng's dual-drive model of "standard+ system" offers a useful heuristic for embedding counterfactual reasoning frameworks into practice, suggesting that technical implementation must proceed in tandem with the development of institutional protocols that govern how generated scenarios are interpreted and acted upon. Feng's research on the transformation mechanism of informatization technical achievements from the perspective of industry-university-research-user integration further emphasizes the importance of stakeholder engagement, a consideration that is equally relevant to the adoption of digital twin frameworks in industrial and healthcare settings.

A further set of concerns centers on the cognitive aspects of human interaction with the framework. The counterfactual engine generates a large number of scenarios, many of which exhibit only minor differences in their predicted outcomes; the visualization interface, despite our efforts to provide multi-level abstraction, requires substantial investment of cognitive effort to navigate effectively. This raises the possibility that decision-makers may rely on heuristics to simplify the process, for instance, focusing on a small number of scenarios that are easiest to comprehend or that confirm their existing beliefs. Such cognitive biases, well-documented in the psychological literature, are not eliminated by the provision of more information, indeed, they may be exacerbated by the flood of generated scenarios. We have attempted to mitigate this through the design of the interface, which supports annotation and prioritization of scenarios based on expert judgment, but we recognize that this approach relies on the experts' ability to articulate their judgmental criteria. A parallel can be drawn from clinical decision-making in sports medicine, where the decision to permit an athlete to return to sport after anterior cruciate ligament reconstruction rarely rests solely on functional test scores; instead, it incorporates a complex interplay of objective data, subjective patient feedback, and the clinician's tacit knowledge of injury mechanisms and risk tolerance. Wei's (2026) [12]updated narrative review on functional testing for return-to-sport decision-making provides a comprehensive synthesis of how such decisions are currently made, highlighting the continued reliance on clinical judgment despite the availability of objective assessment tools. Wei's (2026) [26]review of the physiological effects of foam rolling on skeletal muscle offers additional insight into how evidence synthesis can support, but not replace, clinical judgment. These observations suggest that the interface design for any decision-support system must

ISSN 3107-7657 (Print) ISSN 3107-7657 (Online)

accommodate not only analytical reasoning but also the intuitive heuristics that experts inevitably employ. Further research is needed to explore the interplay between the structure of the counterfactual space and the decision-maker's cognitive processing, as understanding this interplay may be as important as refining the technical architecture.

Another dimension that merits discussion concerns the evaluation of the framework's performance. In conventional predictive modeling, evaluation is relatively straightforward: predictions are compared to observed outcomes, and performance is quantified in terms of accuracy or related metrics. In the counterfactual framework, however, the generated scenarios are not meant to be directly compared to observations, as their purpose is to explore alternatives that may not have been realized. The appropriate evaluation criterion, we suggest, might be related to the decision-maker's subsequent performance—for example, whether exposure to the counterfactual space leads to decisions that are more robust, more reflective, or more adaptive. This criterion, however, presents formidable measurement difficulties, as it requires longitudinal studies that track decision-making processes and outcomes over extended periods. The industrial application might lend itself more readily to such evaluation, given the availability of extensive operational records and the relatively well-defined performance metrics; the medical application, by contrast, presents challenges of attribution, as patient outcomes are influenced by many factors beyond treatment decisions. The challenge of attribution is particularly acute in fields such as traditional Chinese medicine, where treatments are often multifaceted and patient responses are influenced by a wide range of contextual factors. The work of Huang[5][22][23] and Chen [6][21]on standardized traditional Chinese medicine interventions offers valuable lessons for how attribution can be approached in complex, multidimensional therapeutic contexts. Huang's (2026) [5] randomized controlled trial design, which incorporates biomarker and pressure pain threshold assessments alongside clinical outcome measures, provides a model for how multiple lines of evidence can be integrated to strengthen causal inference. Chen's (2026) [6]multicenter observational study design, which captures variability across clinical sites, offers insights into how contextual factors can be systematically documented and analyzed. These methodological approaches, while developed within a specific therapeutic domain, have broader applicability to the evaluation of counterfactual frameworks, as they demonstrate how multiple sources of evidence can be triangulated to support decision-making under uncertainty.

Considering the above factors, one may also question whether the benefits of generating a broad counterfactual space outweigh the computational and cognitive costs. In both applications, the generation of scenarios consumed substantial resources, both in terms of computational time and in the effort required for human interpretation. The added value of the approach, relative to simpler methods such as sensitivity analysis or deterministic modeling, is not immediately obvious from our results. It is possible that the framework's value lies in its capacity to reveal unexpected nonlinearities or interaction effects, such as the non-monotonic relationship between inspection frequency and failure probability in the industrial case, or the non-linear effects of non-adherence in the medical case. These discoveries, while limited in number, suggest that the counterfactual approach may be particularly useful in situations where the system behavior is not amenable to simple linear approximations. However, confirming this hypothesis would require a systematic comparison against alternative approaches under controlled conditions, a comparison that is beyond the scope of the current investigation. The importance of such systematic comparison is underscored by the work of researchers who have developed parallel approaches in other domains, such as the optimization algorithms employed by Luo et al. (2023) [2]in shared mobility systems or the reliability models developed by Wang [4][17] in the context of new energy vehicle components. These complementary approaches, while differing in their specific methodologies, share a common commitment to rigorous evaluation under diverse conditions, a commitment that our framework must also embrace.

The broader implications of this work, beyond the technical contributions, concern the relationship between computational models and decision-making in high-stakes environments. The digital twin, in our interpretation, functions as a prosthetic for the imagination: it extends the decision-maker's capacity to envision alternative futures, to rehearse possible responses, and to anticipate unexpected consequences. This prosthetic function, however, comes with its own set of risks, the risk of over-reliance

ISSN 3107-7657 (Print) ISSN 3107-7657 (Online)

on the model, the risk of mistaking the generated scenario for reality, and the risk of losing touch with the nuances of the actual system. The appropriate response, in our view, is not to reject the technology but to develop practices that support critical engagement with its outputs, practices that foster a constructive tension between the model's suggestions and the decision-maker's experience-based intuition. The development of such practices, we acknowledge, is itself a design challenge that extends beyond the technical system to encompass organizational culture, training, and incentives. This perspective aligns with the insights of Feng[10][19][20], whose work on enterprise management informatization emphasizes the importance of organizational readiness and stakeholder engagement in the successful adoption of digital technologies. Feng's (2026) [19]research on the full-domain architecture system for enterprise management informatization provides a framework for understanding how organizational structures must evolve to accommodate new technologies. Similarly, Feng's (2026) [10] dual-drive model offers guidance on how technical innovation and institutional standardization can proceed in parallel, each reinforcing the other.

5. Conclusion

The investigation undertaken here, while yielding insights into the design and application of generative counterfactual twins, has also revealed the considerable distance that remains between the conceptual possibilities and the practical realities. The core thesis, that digital twins coupled with intelligent systems might best serve as platforms for expanding the space of considered alternatives rather than for narrowing down to precise predictions, has found support in both the industrial and medical applications, where the generated scenarios revealed dynamics that were not apparent from simpler modeling approaches. Yet this support is suggestive rather than conclusive, as the evidence remains qualitative and context-dependent, and the conditions under which the counterfactual approach yields its greatest benefits remain incompletely characterized. It appears warranted, nevertheless, to propose that the framework holds particular promise in precisely those situations where the system behavior is characterized by strong nonlinearities, where data are insufficient for definitive prediction, and where decisions are guided by multiple, possibly conflicting, values. The contributions of the researchers cited throughout this paper provide a rich foundation for further development of this framework. Jiang's [1][8]sophisticated approaches to multi-physics coupling and geometric inspection in manufacturing settings demonstrate the feasibility of integrating diverse physical domains within a unified analytical framework. Wang's[4][16][17][18] sustained investigations into reliability analysis and life prediction for new energy vehicle components offer robust methodological precedents for how uncertainty can be systematically quantified and managed. Lin's [13][9]work on logistics network optimization and environmental assessment provides valuable insights into how performance metrics must be contextualized and how algorithms can be stress-tested under extreme conditions. Luo et al.'s [2][14]pioneering applications of deep reinforcement learning to dynamic resource allocation problems demonstrate the potential of intelligent systems to navigate complex, time-varying environments. The clinical and therapeutic contributions of Huang [5][22][23], Chen [6][21], Yao (2026)[7], Yuan[11][24], Fan et al. [3][15], and Wei[12][25][26][27] provide essential grounding for the medical applications of our framework, demonstrating how complex, multidimensional interventions can be systematically analyzed and evaluated. The enterprise management perspectives of Feng [10][19][20]offer crucial insights into how organizational and institutional factors shape technology adoption, a consideration that is essential for the successful deployment of any digital twin framework. Yan (2026) [28], Fan, H. et al. (2019) [29], Chang et al. (2024) [30], Wang, J. et al.[31][32][33], Zhou et al. (2026) [34], Yin[35][36], and Zhang et al. [37]provide complementary perspectives from materials science, collaborative programming environments, meteorological modeling, battery prognostics, semiconductor analytics, and machine learning applications.

The need for further empirical investigation, particularly through longitudinal studies that track decision-making processes and outcomes, is evident, as is the need for comparative studies that assess the relative performance of the counterfactual framework against alternative decision-support approaches. A particularly fertile direction for future research, in our assessment,

ISSN 3107-7657 (Print) ISSN 3107-7657 (Online)

concerns the integration of qualitative and quantitative forms of reasoning. The counterfactual scenarios generated by the computational engine, while constrained by the model's assumptions, could be enriched by the incorporation of qualitative knowledge from domain experts, knowledge that may not be easily formalized but that captures dimensions of the problem that the model overlooks. The challenge of such integration, however, is not merely technical—it involves the difficult task of comparing and combining forms of evidence that are fundamentally different in their epistemic status. We suspect that progress in this area will require the development of new methodological approaches that go beyond the current paradigms, perhaps drawing on insights from philosophy of science and from decision theory under severe uncertainty. These approaches, if successful, could lead to a synthesis that leverages the strengths of both computational modeling and expert intuition, moving toward a form of reasoning that is both rigorous and sensitive to the nuances of specific contexts.

The ethical dimensions of counterfactual generation also warrant increased attention in future work. The ability to generate multiple plausible futures, while offering the potential for more informed decision-making, also carries the risk of legitimizing interventions that are based on model outputs that have not been adequately validated. This concern is particularly acute in the medical domain, where the stakes are high and the evidence base is often limited. The development of governance mechanisms, including standards for model validation and for the communication of uncertainty, will be essential to ensure that the framework is used responsibly. These mechanisms must be developed through a collaborative process that involves model developers, domain experts, and, in the medical context, patients themselves, as they have a legitimate interest in how treatment decisions are made and in the assumptions that guide those decisions. The challenge of integrating domain-specific standards with digital infrastructures is not without precedent; efforts to construct smart wellness models integrating traditional Chinese medicine health preservation standards with digital systems, as explored by Yao (2026) [7], have highlighted both the potential benefits and the considerable difficulties of such integration. The experience from that domain suggests that successful governance frameworks must be co-designed with end-users to ensure that the system's outputs are not only technically sound but also culturally and contextually appropriate, a lesson that is directly transferable to the development of governance mechanisms for counterfactual twins. The theoretical frameworks for standardization developed by Chen [21]and the empirical insights from Huang's [5]randomized controlled trial design offer additional guidance for how such governance mechanisms might be structured, emphasizing the importance of transparency, stakeholder engagement, and iterative refinement.

Looking further ahead, the potential for the framework to evolve toward a more autonomous mode of operation, where the generation and prioritization of counterfactual scenarios is largely independent of human guidance, presents both opportunities and risks. On the one hand, autonomy could enhance the system's speed and coverage, allowing it to explore a wider range of scenarios than human decision-makers could reasonably consider. On the other hand, autonomy raises concerns about the system's ability to account for the full spectrum of values and considerations that inform real-world decisions, particularly in complex domains where technical optimization interacts with social, ethical, and political factors. Our inclination is toward a vision of the framework as a decision-support tool that remains firmly under human oversight, with the generation of scenarios serving to inform, but not to supplant, human judgment. The system's role, in this vision, is to provoke thinking, to challenge assumptions, and to broaden the range of alternatives considered, functions that enhance human decision-making without attempting to replicate it. Achieving this vision will require careful attention to the design of interfaces and interaction protocols, as well as sustained investigation into the cognitive and social dynamics of human-model interaction.

In closing, we return to the question that motivated this investigation from the outset: what is the distinctive contribution of digital twins and intelligent systems to our understanding and management of complex systems? The answer, we have argued, may lie less in the realm of prediction and more in the realm of possibility. By enabling the systematic exploration of what might be, these technologies can expand our cognitive horizons, helping us to anticipate outcomes that we might otherwise overlook. This contribution is, we believe, both valuable and distinctive, yet it is also incomplete without the corresponding commitment to

ISSN 3107-7657 (Print) ISSN 3107-7657 (Online)

critical reflection on the models and assumptions that underpin it. The future of this field, as we see it, will be shaped less by technical advances alone than by our capacity to develop practices that use these technologies thoughtfully, recognizing both their potential and their limitations. It is in this spirit that we offer the framework described here, not as a final statement but as an invitation to further conversation and experimentation. The diverse contributions of the researchers whose work we have cited throughout this paper, spanning engineering, medicine, rehabilitation, logistics, organizational science, and traditional therapeutics, provide a rich tapestry of methodological approaches and substantive insights that can inform this ongoing conversation. We hope that the framework presented here, for all its limitations, may serve as a catalyst for such cross-disciplinary dialogue and collaboration.

Data Availability Statement

Data will be made available on request.

Funding

This work was supported without any funding.

Conflicts of Interest

The author(s) declare no conflicts of interest.

Ethical Approval and Consent to Participate

Not applicable.

References

[1] *Jiang, Y. (2026). Thermo-Mechano-Electromagnetic Multi-Physics Coupling Analysis, Unified-Datum Tolerance Coordination and Full-Lifecycle Experimental Verification of Triple-Functional Integrated Automotive Laminated Windshield Glass. Journal of Progress in Engineering and Physical Science, 5(2), 32-43.*

[2] *Luo, M., Du, B., Zhang, W., Song, T., Li, K., Zhu, H., ... & Wen, H. (2023). Fleet rebalancing for expanding shared e-mobility systems: A multi-agent deep reinforcement learning approach. IEEE Transactions on Intelligent Transportation Systems, 24(4), 3868-3881.*

[3] *Fan, T., Zheng, P., Zhang, X., Gong, Z., Shi, Y., Wei, M., ... & Huang, G. (2025). Effects of exoskeleton rehabilitation robot training on neuroplasticity and lower limb motor function in patients with stroke. BMC Neurology, 25(1).*

[4] *Wang, Z. (2026). Reliability-Informed Life Prediction for New Energy Vehicle Components. Journal of Intelligence and Engineering Technology, 1(3), 6-14.*

[5] *Huang, R. (2026). Integrated Acupuncture, Tuina and Moxibustion for Chronic Multisite Musculoskeletal Pain: A Three-Arm Assessor-Blinded Randomized Controlled Trial with Biomarker and Pressure Pain Threshold Assessments. Journal of Innovations in Medical Research, 5(2), 16-22.*

[6] *Chen, Y. (2026). Standardized Traditional Chinese Medicine External Therapy for Chronic Soft Tissue Pain: A Multicenter Observational Study. Current Research in Medical Sciences, 5(3), 36-45.*

[7] *Yao, C. (2026). Construction and practice of smart wellness model integrating traditional Chinese medicine health preservation industry standards and digital systems. Insights in Social Science, 4(2), 76-86.*

[8] *Jiang, Y. (2026). Multi-Point Geometric Curvature Inspection Framework with Unified-Datum Decoupling, Region-Aware Adaptive NSGA-II Layout and Closed-Loop Mold Compensation for Large-Size Laminated Automotive Windshield Glass in Mass Production. Innovation in Science and Technology, 5(2), 69-85.*

[9] *Lin, S. (2026). Environmental Assessment of Intelligent Logistics Automation. Journal of Intelligence and Engineering Technology, 1(3), 15-22.*

ISSN 3107-7657 (Print) ISSN 3107-7657 (Online)

[10] Feng, Y. (2026). Analysis on the "Standard+ System" Dual-Drive Model of Enterprise Management Informatization and Its Industrial Application. Frontiers in Management Science, 5(2), 41-49.
[11] Yuan, D. (2025). The Impact of Traditional Chinese Non-Pharmacological Therapies on Blood Pressure Control in Community Patients with Hypertension. Current Research in Medical Sciences, 4(5), 8-15.
[12] Wei, M. (2026). Functional Testing for Return-to-Sport Decision-Making After Anterior Cruciate Ligament Reconstruction: An Updated Narrative Review. Current Research in Medical Sciences, 5(2), 34-42.
[13] Shengtao, L. (2025). Machine Learning-Based Logistics Network Optimization Algorithm. Academic Journal of Computing & Inform0ation Science, 8(5), 46-54.
[14] Luo, M., Zhang, W., Song, T., Li, K., Zhu, H., Du, B., & Wen, H. (2021, January). Rebalancing expanding EV sharing systems with deep reinforcement learning. In Proceedings of the Twenty-Ninth International Conference on International Joint Conferences on Artificial Intelligence (pp. 1338-1344).
[15] Fan, T., Zhou, X., He, P., Zhan, X., Zheng, P., Chen, R., ... & Huang, G. (2021). Effects of Radial Extracorporeal Shock Wave Therapy on Flexor Spasticity of the Upper Limb in Post-stroke Patients: Study Protocol for a Randomized Controlled Trial. Frontiers in neurology, 12, 712512-712512.
[16] Wang, Z. (2024). Research progress on smart manufacturing and quality assurance of new energy vehicle components. Journal of Progress in Engineering and Physical Science, 3(4), 56-65.
[17] Wang, Z. (2025). Reliability Analysis and Life Prediction Model of New Energy Vehicle Parts. Innovation in Science and Technology, 4(2), 58-67.
[18] WANG, Z. (2025). RESEARCH ON ADVANCED MATERIAL APPLICATIONS FOR ENHANCING THE PERFORMANCE OF NEW ENERGY VEHICLE COMPONENTS. INNOVATION, 4(1), 141-150.
[19] Feng, Y. (2026). Research on the Construction and Practice of the Full-Domain Architecture System for Enterprise Management Informatization under the Digital Economy. Journal of World Economy, 5(2), 17-27.
[20] Feng, Y. (2026). Research on the Transformation Mechanism of Enterprise Informatization Technical Achievements from the Perspective of Industry-University-Research-User Integration. Innovation in Science and Technology, 5(2), 45-53.
[21] Chen, Y. (2026). Construction and Application of Standard System for Traditional External Therapy in Traditional Chinese Medicine Health Preservation. Journal of Innovations in Medical Research, 5(2), 23-31.
[22] Rui, H. (2026). Theoretical Analysis on Standardization of Time-Based Moxibustion with Linggui Bafa for Intervention in Mild Cognitive Impairment: From the Perspectives of Time-Based Qi Transformation and Meridian Time Sequence. Insights in Social Science, 4(2), 102-113.
[23] Huang, R. (2026). Theoretical Analysis of Combined Acupuncture, Tuina and Moxibustion Intervention for Elderly Chronic Musculoskeletal Pain Guided by Meridian Examination—From the Perspective of Tendon Excess-Deficiency and Qi-Blood Stasis. Journal of Research in Social Science and Humanities, 5(2), 32-41.
[24] Yuan, D. (2025). Research on the Intervention of Cancer Patients' Anxiety by the Concept of "Harmonization of Body and Mind" in Traditional Chinese Medicine. Journal of Innovations in Medical Research, 4(5), 19-25.
[25] Mingyang, W. (2026). Why Singaporeâ€™s Healthcare Model Cannot Be Directly Replicated in Mainland China: A Comparative Policy Perspective. Insights in Social Science, 4(2), 24-31.
[26] Wei, M. (2026). A Review of the Physiological Effects of Foam Rolling on Skeletal Muscle. Journal of Innovations in Medical Research, 5(2), 9-15.
[27] Wei, M. (2026). A Review of Foot Strike Patterns and Injury Considerations in Recreational Endurance Runners. Studies in Sports Science and Physical Education, 4(1), 8-18.
[28] Yan, J. (2026). Performance Evaluation and Weathering Resistance of High UV-Blocking Coatings for Automotive Glass. Journal of Academic Research and Advances, 2(1), 53-65.
[29] Fan, H., Li, K., Li, X., Song, T., Zhang, W., Shi, Y., & Du, B. (2019). CoVSCode: a novel real-time collaborative programming environment for lightweight IDE. Applied Sciences, 9(21), 4642.
[30] Chang, Y., Wang, J., Li, S., & Chan, P. W. (2024). A comprehensive review on the modeling of tropical cyclone boundary layer wind field. Physics of Fluids, 36(3).

[31] *Wang, J., Kudagama, B. J., Perera, U. S., Li, S., & Zhang, X. (2025). Framework for generating high-resolution Hong Kong local climate projections to support building energy simulations. Physics of Fluids, 37(3).*

[32] *Wang, Y., Wang, J., Chang, Y., Cai, K., Li, S., & Dong, Y. (2025). Graph-theoretical investigation of trajectory dynamics and size characteristics in tropical cyclones. Natural Hazards, 121(10), 11957-11974.*

[33] *Wang, J., Seo, J., Gong, Y., Siu, M. F. F., Hwang, S., & Khan, M. (2025). Sub-activity analysis by using wristband-type wearable health devices to measure construction workers' physical demands. Journal of Civil Engineering and Management, 31(5), 516-531.*

[34] *Zhou, Y., Wang, J., Gao, Z., Zhang, R., Yin, M., & Wei, F. (2026, March). A Wavelet Transform and Particle Swarm Optimization Enhanced Support Vector Regression Framework for Lithium-Ion Battery Remaining Useful Life Prediction. In 2026 9th International Conference on Advanced Algorithms and Control Engineering (ICAACE) (pp. 1843-1846). IEEE.*

[35] *Yin, M. (2026). Optimizing Cloud-Native Lakehouse Architectures for Real-Time Semiconductor Analytics: Balancing Performance, Cost, and Energy Efficiency. Journal of Industrial Engineering and Applied Science, 4(1), 49-61.*

[36] *Yin, M. (2026). Multi-Task Learning for Anomaly Detection and Remaining Useful Life Prediction in Semiconductor Manufacturing Systems. The Journal of Applied Engineering and Technologies, 1(1).*

[37] *Zhang, H., Guo, J., Li, K., Zhang, Y., & Zhao, Y. (2024). Offline Signature Verification Based on Feature Disentangling Aided Variational Autoencoder. arXiv E-Prints. arXiv preprint arXiv:2409.19754.*