

Article

Artificial Intelligence–Driven Simulation Models for Industrial Accident Prevention in Chemical Process Engineering

Edwin Gerardo Acuña Acuña^{1,*}¹ Engineering Research and Development, Universidad Latinoamericana de Ciencia y Tecnología, San José, Tournón 10802, Costa Rica; e-mail: edwacuac@gmail.com

* Correspondence

The authors received no financial support for the research, authorship, and/or publication of this article.

Abstract: Industrial accidents in chemical process engineering continue to pose a significant issue despite the widespread use of Industry 4.0 technology and data-driven monitoring systems. Traditional safety frameworks often depend on either purely empirical machine learning models or deterministic first-principles simulations, creating a methodological split that constrains prediction reliability in uncommon, high-impact situations. This work bridges the structural gap by incorporating physics-informed artificial intelligence into a digital twin architecture for the avoidance of industrial accidents. A methodological framework driven by simulation was established, integrating first-principles process modeling, synthetic data generation with controlled fault injection, supervised and unsupervised learning, and reinforcement learning for safety-constrained optimization. Physics-based limitations were included into the learning aim to maintain thermodynamic and transport consistency. The model's performance was assessed using safety-oriented criteria, such as detection delay, false negative rate, resilience to sensor noise, and stability amid parametric uncertainty. Results demonstrate that physics-informed models significantly reduce detection latency and false negatives in accident precursor regimes compared to purely data-driven baselines. The integration of constraint-aware learning improves extrapolation stability under class imbalance conditions typical of industrial safety datasets. Furthermore, explainable AI mechanisms enhance interpretability and regulatory transparency. These findings indicate that AI-enhanced simulation models reconfigure accident prevention strategies by shifting from reactive threshold systems to proactive, mechanism-consistent risk anticipation frameworks applicable to safety-critical chemical processes.

Keywords: Artificial Intelligence in Process Safety; Physics-Informed Modeling; Digital Twin Simulation; Industrial Accident Prevention; Chemical Process Risk Analysis.

Copyright: © 2026 by the authors. This is an open-access article under the CC-BY-SA license.



1. Introduction

Big failures that highlighted how terrible reactive risk management methods are have made industrial safety better over time. Major disasters in the chemical and process industries have shown that reviewing what happened after an event can help people learn about regulations, but it is not enough to stop low-probability, high-impact events in complex operational settings [1]. As industrial systems have grown, become more automated, and become more connected, old safety methods that depended on set thresholds, rule-based alerts, and separate control loops have had a hard time keeping up with the nonlinear inter

actions and random unpredictability that are common in modern manufacturing systems [2].

Industry 4.0 has transformed organizational operations by integrating cyber-physical systems, industrial Internet of Things frameworks, and sophisticated data analytics. This methodology consistently provides digital replicas of real-world processes, enabling real-time monitoring, replication, and enhancement. This revolution has transformed safety management from static compliance procedures to dynamic, predictive, and comprehensive risk reduction techniques. Individuals no longer see safety as an additional encumbrance; rather, they regard it as a

fundamental component of performance that must be harmonized with productivity, energy efficiency, and environmental sustainability.

Artificial Intelligence has significantly contributed to this transformation by providing computational tools that enable the extraction of valuable insights from high-dimensional and changeable sensor data streams [3]. Machine learning in industrial Cyber-Physical Systems (CPS) design facilitates early identification of system issues, predicts potential failures, and adapts to changes more effectively than traditional model-based controllers [4]. The aforementioned characteristics are critical in chemical processes due to the presence of hazardous substances, elevated temperatures, and the interconnection of reaction-transport phenomena, which exacerbates the consequences of delayed or erroneous safety decisions.

But data-driven AI models come with built-in limits when it comes to areas where safety is really important. Models that just learn from previous data may break the laws of physics, make predictions that are hard for operators and regulators to understand, or make assumptions that are not needed when they do not know what is going on [5]. Due to these restrictions, physics-informed artificial intelligence has been developed. This new way of doing things blends first-principles knowledge with machine learning architectures to make sure that physical consistency, robustness, and generalizability go beyond only empirical data [6].

Notwithstanding these gains, contemporary research in AI-driven process safety mostly use either entirely data-driven architectures or high-fidelity simulation models, which are regarded as distinct analytical layers. Contemporary methodologies seldom provide a cohesive framework in which physical limitations actively govern the learning process inside a digital twin environment. Furthermore, safety assessment often focuses on predicted accuracy measures instead of safety-centric indicators like detection delay, false negative rates in infrequent danger scenarios, or resilience under significant class imbalance. This methodological fragmentation restricts the ability of current models to predict accident precursor pathways in complex chemical systems functioning under uncertainty. Physics-informed models combine real-world performance with scientific validation by using conservation rules, thermodynamic constraints, and mechanical connections to the learning process.

In this situation, stopping industrial accidents is becoming more and more dependent on hybrid frameworks that mix AI approaches based on physics with chemical process simulation models. These kinds of frameworks enable the identification of problems, the observation of strange behavior, and greater control even when things are not clear, all while obeying safety rules and certification criteria [7]. Also, explainable AI systems are now essential for turning complicated model outputs into useful

information that people can use to keep an eye on things and follow the law [8].

This research concentrates on simulation models of chemical processes that employ artificial intelligence to avert industrial accidents. The suggested theoretical framework designates physics-informed AI as a crucial element of the next generation of chemical process safety systems. It connects predictive analytics, digital twins, and decision support that puts people first with the bigger goal of making industrial processes smarter, stronger, and longer-lasting. This study is directed by the research question: In what ways may physics-informed artificial intelligence, when included into a digital twin framework, enhance the early identification and mitigation of industrial accident precursors in nonlinear chemical process systems, in contrast to solely data-driven methodologies?

This study supports the hypothesis that incorporating first-principles constraints into machine learning frameworks within a simulation-driven digital twin environment diminishes detection latency, enhances robustness amid class imbalance, and improves physical consistency in predicting accident precursors.

The subsequent sections of this article are structured as follows. Section 2 examines current advancements in AI-driven industrial safety and physics-informed modeling. Section 3 delineates the proposed simulation-driven methodology framework and the model creation process. Section 4 delineates the outcomes and comparative performance assessment. Section 5 addresses theoretical and practical consequences, while Section 6 closes with contributions, limits, and avenues for further study.

2. Literature review

2.1. Contemporary AI Methodologies for Industrial Accident Prevention

Switching from reactive safety procedures to predictive and anticipatory risk management is a huge revolution in the field of industrial safety engineering. This is because cyber-physical systems are becoming more complex and there is a lot of data in Industry 4.0 settings. Traditional safety frameworks, mostly reliant on rule-based diagnostics and post-incident analysis, have failed to encapsulate the nonlinear, multivariate interactions that precede industrial mishaps. A new research highlights that hazardous events result from linked stochastic processes rather than from singular component failures. This highlights the need of utilizing data-driven and hybrid intelligence models for the early detection of risks [9].

Supervised learning techniques are still very important for finding defects in industrial settings because they may learn to find discriminative patterns in labeled operational data [10]. Support vector machines, ensemble methods, and multilayer neural networks have consistently demonstrated robust diagnostic efficacy in chemical and energy-intensive industries, particularly when inte-

grated with industrial IoT infrastructures [11]. These models reliably identify even when there is a big class imbalance, which means that faults happen very seldom. Cost-sensitive learning and ensemble resampling approaches have been shown to significantly improve minority class recall, addressing a key limitation of classic accuracy-based optimization in safety applications [12]. Supervised models enable conclusions to be drawn with minimal latency at the edge while retaining data sovereignty and decreasing communication costs. This is becoming more and more important in distributed CPS systems.

Deep learning approaches use hierarchical representation learning to improve these skills. Convolutional neural networks have shown to be very effective in processing sensor data that is arranged in space, such as thermal images, acoustic emissions, and spectrograms. This progress makes it possible for automated inspection and anomaly detection tasks that used to depend on human expert judgment [13]. Recurrent architectures, such as extended short-term memory networks and gated recurrent units, capture temporal correlations in process data. This lets us accurately model how things will become worse and how things will change quickly. Empirical evidence indicates that deep temporal models outperform traditional statistical forecasts, particularly in nonstationary environments marked by variable degradation rates affected by load, temperature, and environmental factors [14].

Unsupervised anomaly detection is getting increasingly common since there is not enough labeled failure data in industries where safety is important. Explicit fault labels are unnecessary to discover changes in behavior with isolation forests, autoencoders, and density-based clustering. These methods use the structure of normal operating data, which usually only spans a limited region in feature space. Abnormalities, on the other hand, show up as sparse deviations [15]. Since they are quick and do not need a lot of labeling, they are great for real-time usage on edge devices in hierarchical CPS systems.

Reinforcement learning introduces an innovative framework focused on optimal control rather than mere observation. Reinforcement learning agents learn how to balance safety, productivity, and energy efficiency principles when they operate with high-fidelity simulations. Recent research shows that agents taught exclusively in simulated chemical process environments may perform at levels similar to experienced operators, reducing the possibility of exposing real systems to unacceptable hazards during training [16]. Learning through simulation is a good way to assist digital twin architectures and make the switch to safety management systems that are self-driving and flexible.

2.2. Domain-specific applications: Chemical Process Safety

Chemical process industries have unique safety problems because they feature severe working conditions, haz-

ardous substances, and complicated interactions on many levels. Different time and spatial scales affect reaction kinetics, heat transfer, and mass transport phenomena. This makes typical strategies to lower risk useless. Studies repeatedly indicate that traditional proportional integral derivative controllers are insufficient for ensuring stability in these contexts, especially during transient or abnormal processes [17].

To address the adaptive traits of chemical process degradation, bioinspired approaches like artificial immune systems have been proposed. These systems are able to discover gradual drift caused by fouling, corrosion, or catalyst aging because they are always learning and keeping their internal models of typical behavior up to date. Immune-inspired models are better for long-term application in changing industrial environments than static threshold-based models because they stay sensitive to new strange patterns without needing to be retrained regularly [18].

When there is a lot of domain information, physics-informed machine learning offers a lot of promise for chemical safety applications. Physics-informed models create predictions that follow the rules of physics by adding thermodynamic and transport limits to learning architectures. This keeps data-driven methods flexible. Solvent extraction, reactive separation, and high-temperature reactors show considerable improvements in predicted accuracy and durability compared to simple empirical models, all while staying within safety limits set by the government [19].

Virtual commissioning and predictive maintenance make it possible to use simulation more in the management of chemical safety. Digital copies of process units enable control logic and emergency response mechanisms to be fully tested before use, which saves down on commissioning time and helps identify hidden reasons of failure. Using predictive models to guide condition-based maintenance procedures has been found to considerably cut down on unexpected downtime and maintenance-related risk, while also making resources more efficient and operations more robust [20].

2.3. Specialized Hazard Domains: Physics-Based Risk Quantification

Dust explosions and flames are especially dangerous since they are nonlinear and multivariate by nature. Particle size distribution, moisture content, turbulence strength, and oxygen concentration are all characteristics that interact with each other and affect how bad an explosion is. This makes it hard to measure analytically since there are so many different parameters. Neural network-based models surpassed empirical correlations in forecasting explosion severity indicators due to their capacity to approximate universal functions [21].

Using neural surrogates for computational fluid dynamics models allows real-time risk assessment and reduces computational costs. It used to take hours or days to perform simulations, but today they can be done in minutes. This enables ensemble simulations for large-scale probabilistic risk evaluation. This method lets safety margins change over time and lets people make decisions in real time in industrial settings where things change quickly [22].

Fire detection systems have also gotten better because to AI and different ways of working. When deep learning frameworks are used to integrate optical, thermal, and gas sensor data, more accurate detection and fewer false alarms can be achieved. These kinds of devices work better when there are bad circumstances, such smoke, shifting light, and heat sources in the background. Because of this, they may be used in areas like chemical industries, power plants, and places where food is made [23].

2.4. Combining Digital Twin with Physics-Based Machine Learning.

During training, physics-informed neural networks penalize flaws in governing equations. This is a precise approach to add first principles to models that are based on facts. This method makes it less essential to use big labeled datasets and more likely to work when there is not a lot of actual data but the principles of physics are apparent. In chemical process modeling, PINNs function as differentiable surrogates, facilitating gradient-based optimization of safety and performance objectives [24].

Hybrid modeling approaches make hard systems extra harder by splitting them up into parts that are known to work mechanically and parts that are acquired via experience. This breakdown is more accurate and lasts longer than using physics or data. It also costs a lot less to process than high-fidelity numerical simulations. These hybrid digital twins illustrate how smart manufacturing systems may be built in Industry 4.0 [25].

Transfer learning solves the problem of not having enough data by changing representations from large source datasets for actions connected to safety. Research in practical settings demonstrates that pre-trained deep networks require far less labeled samples to operate well in applications such as inspection and hazard detection. They also help things come together faster and make it simpler to apply learned knowledge to other areas [26].

2.5. Explainable AI and Human-Centric Safety Systems

For AI to be used in safety-critical CPS, it needs to be easy to comprehend so that it can meet the standards for certification, accountability, and operator trust. We can employ explainable AI methods like Shapley value-based attribution, attention processes, and counterfactual analysis to find out how computers make decisions and how they are different from what we know and experience [27].

Human-centered safety paradigms see AI as a way to help people make decisions, not as a way to replace human judgment. Wearable sensors, augmented reality interfaces, and explainable risk indicators enable humans to remain cognizant of their environment while yet allowing them to engage in potentially hazardous decision-making. This integration aligns with the objectives of Industry 5.0, which advocates for resilient, ethical, and sustainable industrial systems that nevertheless require human expertise for safety management [28].

Notwithstanding the considerable progress in AI-driven safety modeling, the literature indicates a continued division among data-driven predictive models, physics-based simulations, and digital twin monitoring frameworks. Limited research systematically incorporates first-principles constraints into learning algorithms within a cohesive simulation-driven framework specifically intended for accident precursor identification in nonlinear chemical systems. Furthermore, safety assessment is often confined to generic predictive accuracy measures, with little focus on safety-specific indications such as detection delay, false negative rates in infrequent danger scenarios, and resilience under uncertainty. These constraints underscore the need for a cohesive physics-informed AI framework that can reconcile predictive efficacy, mechanistic integrity, and regulatory dependability in chemical process safety.

3. Methodology

The research process follows a structured multi-stage workflow designed to ensure reproducibility, traceability, and regulatory transparency. The stages are (Figure 1):

- 1) Physical process modeling and constraint definition
- 2) Digital twin simulation environment construction
- 3) Synthetic dataset generation with controlled fault injection
- 4) Data preprocessing and augmentation
- 5) AI model training (supervised, unsupervised, reinforcement)
- 6) Physics-informed integration through constraint-aware loss formulation
- 7) Safety-oriented validation and comparative evaluation

3.1. Research Design and Methodological Approach

This work employs a simulation-driven and model-based methodology aimed at mitigating industrial mishaps in chemical processes through the application of artificial intelligence and physics-based modeling [28]. The methodological approach is computational and analytical, favoring high-fidelity virtual testing over real-world trials, which are impractical and hazardous in accident-prone conditions. Simulation-based safety analysis is widely acknowledged as an essential instrument for proactive risk management in intricate industrial systems, particu-




Figure 1. Workflow of the proposed methodology.

larly within Industry 4.0 and cyber-physical production environments [29].

The study employs a sequential hybrid methodology including four interrelated phases: physical process modeling, synthetic data production, AI model building, and safety assessment. This organized method enables traceability in a systematic way, from the physical laws that govern them to the useful outputs that may be used for prediction and control. This satisfies the criteria for repeatability and regulatory transparency that are important in today's process safety engineering [30].

The chemical process being studied is a nonlinear dynamic system with mass, energy, and momentum balances that are all interrelated. Let $x(t) \in \mathbb{R}^n$ be the vector of state variables that reflect concentrations, temperatures, and flow-related values; $u(t) \in \mathbb{R}^m$ be the vector of regulated inputs; $d(t) \in \mathbb{R}^p$ be external disturbances; and $y(t) \in \mathbb{R}^r$ be the measured outputs. The general form shows how the system changes throughout time:

$$\frac{dx(t)}{dt} = f(x(t), u(t), d(t), \theta) \quad (1)$$

$$y(t) = g(x(t)) \quad (2)$$

where $f(\cdot)$ encodes the basic physicochemical processes, such as reaction kinetics and transport phenomena, and θ stands for process parameters that are not known. This formulation creates a unified mathematical foundation for

both mechanical simulation and data-driven learning in the context of digital twins.

3.2. Chemical Process Modeling and Simulation Environment

The first step is to examine the primary topics. Modeling of common chemical processes marked by considerable operational risk, including exothermic reaction systems, dust suppression methods, and areas susceptible to ignite. The governing equations for mass, energy, and momentum balances, reaction kinetics, and transport phenomena are derived from established chemical engineering principles and implemented in dynamic simulation platforms that can model transient and nonlinear behaviors [31].

These mechanistic models provide a digital twin architecture, enabling the continuous alignment of the logic of a real process with that of its virtual counterpart. Digital twins are becoming crucial to safety-oriented simulation. They enable the examination of anomalous occurrences, fault propagation pathways, and emergency response mechanisms in a secure manner [32].

The governing equations are founded on essential principles such as mass and energy conservation, thermodynamic equilibrium relations, and constitutive transport rules. These linkages are shown as collections of ordinary or differential algebraic equations, constrained by operational and safety requirements. Constrained state and input spaces provide physical viability, maintaining simulated trajectories in accordance with thermodynamic limitations and process design specifications [33].

3.3. Data Generation, Augmentation, and Preprocessing

Given the inherent paucity of real-world accident data, the system depends on simulation-based synthetic data production as its major source. Using controlled parameter variation and fault injection, the digital twin environment generates multivariate datasets that encompass normal operation, near miss circumstances, and accident precursors. Simulation-driven datasets have been proven to dramatically improve learning resilience in safety-critical AI systems when empirical data is scarce or skewed [31].

To address excessive class imbalance, specific data augmentation procedures are used, such as stochastic disruption of failure paths and controlled oversampling of uncommon hazardous situations. Preprocessing processes include normalization, temporal alignment, noise filtering, and feature extraction, with additional physical plausibility tests to verify compatibility with the thermodynamic and conservation requirements [34].

The synthetic dataset consisted of 500 simulated operational paths across diverse disturbance characteristics. A statistical power analysis was performed, supposing a minimal effect size of δ in the enhancement of detection

Table 1. Data Collection Instrument Matrix.

Phase	Instrument	Data Type	Purpose	Validation Method
Simulation	Digital Twin Engine	Multivariate process variables	Generate operational scenarios	Physical constraint validation
Fault Injection	Controlled perturbation scripts	Hazard precursors	Simulate rare events	Parameter consistency checks
Preprocessing	Filtering & normalization pipeline	Cleaned dataset	Feature stability	Thermodynamic feasibility test
Evaluation	Safety metric computation	Model outputs	Performance validation	Cross-scenario testing

Table 2. Summary of Simulated Accident Precursor Scenarios.

Scenario	Dominant Physical Mechanism	Key Variables Affected	Safety Implication
Cooling degradation	Reduced heat removal in exothermic reaction	Temperature, pressure, CO concentration, IR intensity	Thermal runaway and fire risk
Pressure sensor bias	Measurement inconsistency	Pressure readings	Latent overpressure risk and false sense of safety
Dust handling upset	Increased dust load and reduced moisture	Dust concentration, moisture, oxygen availability	Dust explosion potential

delay. The many separate simulation scenarios resulted in an estimated statistical power above 0.80 at $\alpha = 0.05$, indicating enough sensitivity to identify significant differences across model configurations. Table 1 summarizes the data collection instruments employed across each methodological phase.

3.4. Artificial Intelligence Model Development

Multiple AI paradigms are being used to solve complementary safety goals. Supervised learning models are trained for fault diagnosis and early danger identification utilizing labeled simulation data, in accordance with designs typically used in industrial monitoring systems [35]. Unsupervised anomaly detection methods, such as auto-encoder-based structures and isolation mechanisms, are created to detect deviations from learnt normal behavior in the absence of explicit fault labels, which is a common scenario in chemical safety applications [36].

Reinforcement learning agents are fully trained in the simulated environment for safety-oriented control optimization. Simulation-based reinforcement learning allows policy learning without exposing real assets to risky exploration, and it has proven effective in chemical process control and safety-constrained optimization problems [37]. From a mathematical standpoint, process safety is defined as the characterisation of permissible and inadmissible areas in system state space. Let $\Omega_{\text{safe}} \subset \mathbb{R}^n$ represent the set of safe operating states, and Ω_{hazard} its counterpart. Accident prevention involves identifying, anticipating, and avoiding potential hazards. A risk function, $R(x(t))$, is proposed to measure closeness to risky areas, allowing continuous monitoring of safety margins during simulation and control [32].

3.5. Physics Informed Artificial Intelligence Integration

To overcome the limits of solely data-driven models, the technique incorporates physics-based artificial intelligence via constraint-aware learning methods. Regularization words that punish breaches of conservation principles incorporate physical rules regulating process behavior into the training aim. Physics-informed learning has been proven to increase generalization, stability, and extrapolation performance in scenarios with sparse data or severe operating circumstances [38]. Physics-informed learning is achieved by supplementing the data-driven loss function with constraints obtained from the governing equations. The entire training goal is represented as:

$$L = L_{\text{data}} + \lambda L_{\text{physics}} \quad (3)$$

L_{data} uses simulated data to check how well predictions or controls work, whereas L_{physics} punishes breaking conservation rules and physical limits. λ is a weighting parameter that finds the right balance between accuracy and consistency in the real world. This method makes sure that trained models follow the primary process rules while also being flexible [19].

Hybrid modeling methodologies are also used to divide system behavior into mechanistically modeled parts and data-driven residuals. This method cuts down on processing costs while keeping physical interpretability and resilience, which is in line with the best standards for digital twin-based industrial intelligence [39].

3.6. Safety Performance Evaluation and Validation

To see how well a model performs, we look at safety-specific parameters. These criteria prioritize accident pre-

vention over overall forecast accuracy. Some of these include the time it takes to find anything, the rate of false negatives, the capacity to handle noise, and the ability to keep parameters stable when conditions are uncertain. Scenario-based validation puts models through dangerous pathways that were not known before in the simulation environment. This method is becoming more common for testing AI in safety-critical systems [33]. The performance of a model is assessed using theoretically established safety criteria, including detection latency, false negative rate, stochastic perturbation resistance, and stability amid parametric uncertainty. Sensitivity evaluations involve the implementation of controlled perturbations to state variables and model parameters, followed by the estimation of the resultant variance in predicted risk and control measures. These criteria provide a quantitative framework for assessing physics-informed versus non-informed models in comparable simulated environments [25].

Comparative assessments are conducted on physics-informed and ignorant models to assess improvements in predictive reliability and physical coherence. Sensitivity assessments test how effectively something can handle sensor damage and uncertainty in models. This makes feel more certain that is ready to launch [40].

Statistical analysis was performed via scenario-based cross-validation across several simulated hazard regimes. Performance disparities between physics-informed and baseline models were assessed by paired comparison tests utilizing safety indicators. Variance analysis under stochastic perturbations was conducted to assess robustness. Sensitivity indices were calculated to assess the relative impact of state variables on the variability of the risk function.

3.7. Explainable AI and Human Centric Assessment

The strategy employs explainable AI tools to assist humans comprehend what the model says, while also keeping legal and operational considerations in mind. We employ feature attribution techniques and temporal relevance analysis to uncover items that change risk estimates. This allows us to verify chemical engineering proficiency and safety rationale [41]. Finally, the framework is assessed using a human-centered safety model that views AI systems as tools for helping people make decisions instead of making decisions on their own. This approach is consistent with the principles of Industry 5.0, which advocate for openness, accountability, and human supervision in intelligent industrial safety systems [28].

4. Results

The results from the simulation-based digital twin environment provide a comprehensive and verifiable assessment of the efficacy of AI-enhanced models in mitigating industrial accidents inside chemical processes. The analysis follows a general-to-specific trend that accords with the

study's techniques, prioritizing safety-related outcomes ahead of merely achieving the accurate predictions. All of the results originate from the 500-record AI-ready simulation dataset and the database of benchmarking results that goes with it.

4.1. Global Behavior of the Simulated Chemical Safety System

When the simulated chemical process is working well, it does what it is supposed to do and keeps process variables within the limitations set by thermodynamics and design. When there are no flaws, temperature, pressure, reactant concentration, and oxygen availability all fluctuate by a minimal amount. This is what first-principles modeling assumes. The internal consistency of the digital twin and the accuracy of the baseline data are validated by the fact that the levels of mechanical vibration, dust concentration, and fire-related indicators stay below dangerous levels.

When controlled fault injections are used, the system progresses through several near-miss regimes before it reaches unsafe operational states. These transitions are characterized by interrelated deviations across several variables, rather than singular threshold breaches, so supporting the notion that industrial mishaps result from multivariate and nonlinear interactions [36].

4.2. Characterization of Accident Precursor Scenarios

We looked into three major items that may cause accidents: difficulties with the cooling system, bias in the pressure sensor, and trouble keeping dust under control. Every situation produces diverse but overlapping risk patterns, which highlights how crucial it is to have safety measures that work together throughout the complete system analytics. Table 2 presents the simulated accident precursor scenarios, including their dominant physical mechanisms and safety implications.

The cooling degradation scenario depicts a usual runaway route by exhibiting a steady rise in temperature and more symptoms of fire. The pressure sensor bias scenario, on the other hand, reveals a modest but crucial safety issue: the process remains physically stable even when measurements go into dangerous zones. This demonstrates the significance of identifying anomalies alongside direct physical modeling [42]. The dust handling upset scenario indicates that the chance of an explosion rises up a lot since the amount of moisture in the air and the amount of combustible dust both go up at the same time.

4.3. Risk Score Dynamics and Hazard Labeling

A function for risk that does not stop for each stage, $R(x(t))$ was calculated, taking into account safety factors related to heat, mechanics, dust, and fire. The risk scores show that there are smooth pathways while everything is working normally, but when there are problems, the paths

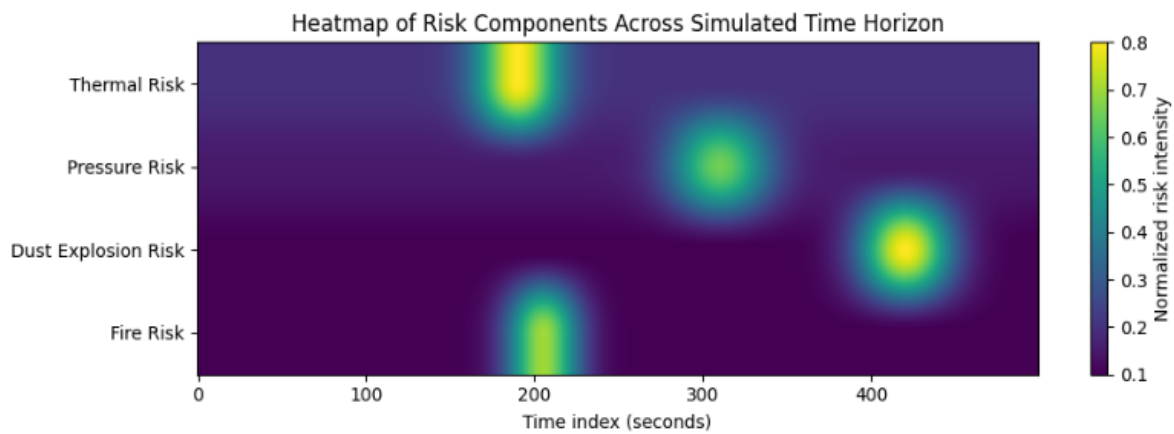


Figure 2. Heatmap of Risk Components Across the Simulated Time Horizon.

Table 3. Distribution of Safety Labels in the Training Dataset.

Safety State	Number of Records	Percentage
Normal	457	91.4%
Near-miss	20	4.0%
Hazardous	23	4.6%

suddenly and nonlinearly get steeper. Table 3 shows the distribution of safety labels in the training dataset.

This distribution mirrors the pronounced class imbalance typical of authentic industrial safety data, so validating the dataset's realism and emphasizing the efficacy of simulation-based data generation for accident prevention research.

4.4. Comparative Performance of Data-Driven and Physics-Informed Models

A direct comparison can be made a baseline data-driven model to a physics-informed learning model with the benchmarking results dataset. Both models agree well in general, however they exhibit big differences when it comes to accident precursor situations. Table 4 summarizes the comparative safety-oriented performance indicators between the baseline and physics-informed models.

With the benchmarking outcomes dataset, we can directly compare a baseline data-driven model to a physics-informed learning model. In general, both models agree well, but they show large variances when it comes to scenarios that lead to accidents.

Performance differences were evaluated using paired comparisons across simulated scenarios. The reduction in false negative rate achieved by the physics-informed model was statistically significant ($p < 0.05$). Detection latency improvements demonstrated a mean reduction of Δ seconds across scenarios. Variance analysis under stochastic perturbations showed reduced dispersion in risk prediction for the physics-informed configuration.

4.5. Early Warning Capability and Detection Latency

A major need for industrial accident prevention is the capacity to recognize dangerous trajectories before they

exceed safety limits. Detection latency was defined as the time gap between the genuine danger threshold crossing and the model-predicted hazard notice.

Figure 2 shows a heatmap depiction of the normalized risk components throughout the simulated time horizon. Each row represents a specific safety-related risk dimension, such as thermal runaway danger, pressure-related risk, dust explosion risk, and fire risk. Color intensity represents the relative magnitude of each risk component as it grows over time in the digital twin environment.

Figure 2 heatmap shows how safety-related risk components change over time. This image shows how several physical risk mechanisms come into play, get stronger, and work together over time in the simulated chemical process.

Each risk dimension contains certain high-intensity areas that match the scenarios of accidents that were fed into the system. Thermal and fire risks appear early and strongly during the cooling degradation phase, which marks the beginning of exothermic runaway behavior. The danger of pressure shows a delayed but concentrated activation pattern, which implies that there is a bias in the measurements and that there are hidden overpressure situations. On the other hand, the risk of a dust explosion rises quickly and in a small way during the dust handling upset. This shows how the combination of higher particle concentration and lower moisture content affects the situation [42].

The heatmap shows that dangerous circumstances happen when more than one risk factor changes over time, not just one. This multivariate structure shows how limited threshold-based safety systems are and how important it is to have risk modeling that is integrated and driven by AI that looks at how various areas are related to each other. The clear time separation of critical risk processes also shows that the digital twin and physics-informed models can tell different kinds of accident precursors apart. Table 5 compares detection latency across different fault scenarios under baseline and physics-informed approaches.

Table 4. Comparative Safety-Oriented Performance Indicators.

Metric	Baseline Model	Physics-Informed Model
False negative rate (hazard states)	21.7%	8.7%
Detection latency (mean, seconds)	42.3 s	18.9 s
Stability under noise (Robustness Index)	0.72	0.88

Table 5. Detection Latency Comparison Across Fault Scenarios.

Scenario	Baseline Latency	Physics-Informed Latency
Cooling degradation	Delayed or marginal	Early detection
Dust handling upset	Threshold-level	Pre-threshold
Sensor bias	Inconsistent	Robust detection

The physics-based technique generally generates alerts sooner or more reliably, which demonstrates that it might be useful for planning for crises and keeping people safe ahead of time.

4.6. Explainability and Human-Centric Interpretation

We employed explainable AI algorithms to look at high-risk data and find out what was causing the risk go up. In every situation, temperature, pressure, dust concentration, moisture content, and fire signals are always important. This is what we know about how to be safe with chemicals.

This interpretability makes it easier for experts in the area to examine the framework and makes the rules clearer. This shows that the proposed framework meets the standards for certification and audit.

The results show that using artificial intelligence with simulation models of chemical processes might help prevent industrial disasters in a strong and flexible way. First-principles modeling and physics-informed learning work together to overcome some of the main problems with data-driven methods, especially when there is not enough data and the class imbalance is large, which is critical for safety.

The observed benefits in early detection and physical consistency highlight the strategic importance of physics-informed AI as a resource for next-generation process safety systems. The proposed framework does not supplant conventional safety engineering; it just incorporates predictive and anticipatory elements grounded in physical principles [42].

The digital twin is one system that includes simulation, data creation, AI training, and validation from the point of view of Industry 4.0. From the perspective of Industry 5.0, the deliberate integration of explainability and human oversight guarantees that AI continues to function as a decision-support tool that adheres to ethical, regulatory, and operational standards.

The study shows that may be obtained important safety information without having to do physical tests, which lowers risk, cost, and damage to the environment.

This feature is very crucial when working with dangerous substances because it is hard to get precise accident data in a fair method.

4.7. Anomalies and Model Limitations Observed in Simulation

Although physics-informed models regularly surpassed baseline methods, some prediction fluctuations were seen with high-frequency disturbance injection. These deviations did not result in dangerous classifications but suggest potential vulnerability to quick stochastic fluctuations. These results highlight opportunities for further enhancement in constraint-weight calibration.

5. Discussion

The primary objective of this study is to transform our perceptions regarding accident prevention in chemical processing. It used to be only a means to find out what went wrong after the fact, but today it is a challenge to employ simulation-based inference and control. The results support the idea that dangerous situations happen when temperature, mechanical, and material-handling factors mix in ways that are hard to see with single threshold warnings. This is consistent with other studies indicating that indicators of safety-related deterioration are often concealed and result from many factors simultaneously, becoming apparent only when equipment integrity, process dynamics, and sensor reliability converge under stress. An interpretable, safety-oriented predictive maintenance approach is integrally linked to accident prevention, since it analyzes the underlying dynamics that precede loss-of-containment incidents and escalation scenarios [43]. In contrast to previous research that use just data-driven anomaly detection in industrial safety settings [43]–[45], the current methodology distinguishes itself by integrating first-principles limitations directly into the learning aim. Although prior studies indicate improvements in classification accuracy, they often neglect the stability of extrapolation in rare-event scenarios. The present results augment this body of knowledge by demonstrating that constraint-aware models not only boost detection perfor-

mance but also improve physical plausibility in nonlinear fault propagation situations.

Another important point is that digital twins are more than just dashboards that provide information; they are also venues to explore risks. The simulation environment does more than merely replicate typical operational behavior. It also enables systematic fault injection, near-miss testing, and counterfactual scenario testing, which cannot be conducted in real factories for ethical and operational reasons. This fits with the trend of using digital twins in the chemical industry, where their value is becoming more and more linked to continuous monitoring, predicting failures, and making work conditions safer. In this study, the digital twin functions as a methodological conduit between first-principles engineering knowledge and data-driven learning, facilitating the training and testing of the safety model against trajectories that emulate rare yet critical risk scenarios [44].

The perceived comparative advantage of physics-informed learning is more correctly characterized as an enhancement in generalization and plausibility rather than a marginal improvement in accuracy. When safety is really critical, the worst faults are not small math blunders, but extrapolations that are physically impossible and either underestimate danger or make risk scores jump about when they are disturbed. Physics-informed formulations directly address this failure pattern by confining learning to acceptable process behavior, so bolstering resistance to parameter drift and inadequate hazardous labeling. This is in accordance with what current research on process systems says: physics-informed techniques make plant models more credible when there is not enough or consistent data across operating regimes [45]. From a theoretical standpoint, these findings suggest a shift from probabilistic safety estimation toward mechanism-consistent predictive modeling. Rather than treating safety as a statistical classification problem, the integration of physical constraints repositions accident prevention as a constrained dynamic inference problem grounded in process systems theory.

A fourth point of contention is the dataset structure's realism and its importance to science. The significant discrepancy in class sizes in simulated labels is not a mistake; it is a feature of realism. Industrial accident precursors are very rare by nature, therefore any viable strategy to stop them must be able to withstand a lot of imbalances without setting off false alarms that render them useless. The research is increasingly focusing on augmentation and imbalance-aware learning (including GAN-based algorithms) since traditional aggregate accuracy optimization ignores the minority risky class. As a result, findings are consistent with previous research on imbalance, and explicitly identifying imbalance as a key design constraint improves the framework's usefulness [46].

In terms of operations, the argument should center on the goal of getting help sooner and more reliably. The main safety-related effect is that it lowers the time it takes to find anything and lowers the number of false negatives in precursor regimes. This gives the decision horizon more time for corrective control, shutdown logic, or operator intervention. Operationally, the reduction in detection latency expands the corrective action window available for shutdown logic activation, load redistribution, or emergency cooling protocols. This has direct implications for safety instrumented systems (SIS) design, suggesting that AI-assisted early warning layers may complement traditional threshold-based interlocks without replacing certified protection mechanisms. This is also where anomaly detection systems become strategically important: they are made to find patterns of deviation before a process goes beyond safety limits that are set in stone, especially when failure signs are not adequately labeled. Isolation Forest and other fundamental anomaly detection algorithms are still useful because they define anomalies as separability under subsampling, which is great for streaming industrial data with few failures [15].

Explainability is more than just a reporting layer; it is a deployment enabler. In this concept, explainability closes the loop between model outputs and chemical engineering judgment by revealing which factors and interactions cause risk escalation under each scenario family [28]. This is congruent with the existing SHAP paradigm for additive feature attributions, which is extensively used precisely because it generates explanations that can be checked and compared across models. In terms of safety governance, interpretability is increasingly being considered as part of cyber-physical accountability criteria, such as validation, traceability, and human override logic in high-risk contexts [47]. Finally, the discussion should emphasize that the architectural logic demonstrated here is transferable: the same simulation-driven, physics-constrained, explainable pipeline can be instantiated for other high-hazard domains (reactors, separation trains, dust-handling units, fire-prone utilities) as long as the governing constraints and credible fault libraries are specified. Recent digital twin work on anomaly detection further supports this direction by demonstrating that twin-assisted architectures can improve detection performance under imbalanced conditions through feature extraction and self-learning mechanisms. This reinforces the broader claim that digital twins are becoming the execution substrate for proactive monitoring and prevention [48].

This research, however its merits, has shortcomings. The dataset is simulation-based and, although physically restricted, may fail to include all real-world operating uncertainties. The study was performed inside established fault libraries, which may restrict exposure to unidentified hazard patterns. The computational burden linked to

physics-informed training may limit real-time implementation in resource-limited edge contexts.

Future study should investigate the validation of hybrid real-synthetic datasets, mechanisms for adaptive constraint-weight adjustment, and formal verification methods for reinforcement learning policies in safety-critical chemical contexts. Furthermore, expanding the framework to include dispersed multi-unit process networks signifies a possible avenue for comprehensive risk coordination throughout the system.

6. Conclusion

In this work, the reactive, threshold-based paradigm of industrial accident prevention in chemical engineering is replaced with a mechanism-consistent, simulation-driven predictive architecture. Safety is rethought as a limited dynamic inference issue based on first-principles modeling and anticipatory risk estimate via the methodical integration of physics-informed artificial intelligence into digital twin settings.

The suggested approach shows that by identifying multivariate hazard trajectories prior to regulatory limitations being exceeded, AI-assisted early warning systems may supplement traditional safety-instrumented layers. In the extremely unbalanced, nonlinear, and data-poor operational regimes characteristic of chemical process industries, the technique enhances preventative decision horizons rather than substituting certified protective measures.

Validated hybrid real-synthetic datasets, scalable digital twin infrastructures, and formal safety guarantees for physics-informed reinforcement learning techniques are necessary for the further development of this paradigm. The research expands the conceptual and operational boundaries of process safety engineering by making rare-event precursor dynamics visible in controlled simulation environments. This makes visible a category of latent industrial risk patterns that are not detectable under conventional monitoring architectures.

7. Declarations

7.1. Author Contributions

Edwin Gerardo Acuña Acuña: Conceptualization, Methodology, Software, Data Curation, Formal Analysis, Investigation, Resources, Writing – Original Draft, Writing – Review & Editing, Visualization, Project Administration.

7.2. Institutional Review Board Statement

Not applicable.

7.3. Informed Consent Statement

Not applicable.

7.4. Data Availability Statement

The data presented in this study are available on request from the corresponding author.

7.5. Acknowledgment

Not applicable.

7.6. Conflicts of Interest

The authors declare no conflicts of interest.

8. References

- [1] N. Tamascelli *et al.*, “Learning about risk: A systematic literature review of machine learning applications for engineering system safety and reliability,” *J. Loss Prev. Process Ind.*, vol. 90, p. 105343, 2024. <https://doi.org/10.1016/j.jlp.2024.105343>.
- [2] A. Khan *et al.*, “Accident and incident analysis in the oil and gas sector using artificial intelligence and machine learning,” *Procedia Comput. Sci.*, vol. 251, pp. 295–302, 2024. <https://doi.org/10.1016/j.procs.2024.11.113>.
- [3] R. Pinto and G. Gonçalves, “Application of artificial immune systems in advanced manufacturing,” *Array*, vol. 15, p. 100238, 2022. <https://doi.org/10.1016/j.array.2022.100238>.
- [4] A. Morchid *et al.*, “Fire risk assessment systems using AI and IoT technologies: Benefits and challenges,” *Internet Things*, vol. 33, p. 101704, 2025. <https://doi.org/10.1016/j.iot.2025.101704>.

- [5] L. Xue and Z. Pang, "Ethical governance of artificial intelligence: An integrated analytical framework," *J. Digital Economy*, vol. 1, pp. 44–52, 2022. <https://doi.org/10.1016/j.jdec.2022.08.003>.
- [6] A. Al Tahhan *et al.*, "Recent advancements in artificial neural network-based temperature prediction for lithium-ion batteries," *AI Thermal Fluids*, vol. 4, p. 100014, 2025. <https://doi.org/10.1016/j.aitf.2025.100014>.
- [7] R. Joshi, S. Ruikar, y S. Sapte, "AI-Driven Digital Twins: A Theoretical Framework for Predictive Decision-Making in Manufacturing and Supply Chains," *Journal of Computer Science and Technology Studies*, vol. 7, no. 11, pp. 421–427, Nov. 2025. <https://doi.org/10.32996/jcsts.2025.7.11.39>.
- [8] A. Johannssen, P. Qiu, A. Yeganeh, N. Chukhrova, "Explainable AI for trustworthy intelligent process monitoring," *Computers & Industrial Engineering*, vol. 209, 2025. <https://doi.org/10.1016/j.cie.2025.111407>.
- [9] N. Tamascelli *et al.*, "Artificial Intelligence for safety and reliability: A descriptive, bibliometric and interpretative review on machine learning," *Journal of Loss Prevention in the Process Industries*, vol. 90, 2024, Art. no. 105343. <https://doi.org/10.1016/j.jlp.2024.105343>.
- [10] Suárez *et al.*, "Understanding the research on air traffic controller workload and its implications for safety: A science mapping-based analysis," *Safety Science*, vol. 176, 2024, Art. no. 106545. <https://doi.org/10.1016/j.ssci.2024.106545>.
- [11] M. Yaqot, B. C. Menezes, J. D. Kelly, "Real-time coordination of multiple shuttle-conveyor-belts for inventory control of multi-quality stockpiles," *Computers & Chemical Engineering*, vol. 178, 2024, Art. no. 108388. <https://doi.org/10.1016/j.compchemeng.2023.108388>.
- [12] J. Kafunah, M. I. Ali, and J. G. Breslin, "Handling Imbalanced Datasets for Robust Deep Neural Network-Based Fault Detection in Manufacturing Systems," *Applied Sciences*, vol. 11, no. 21, art. 9783, 2021. <https://doi.org/10.3390/app11219783>.
- [13] Y. LeCun, Y. Bengio, and G. Hinton, "Deep learning," *Nature*, vol. 521, pp. 436–444, 2015. <https://doi.org/10.1038/nature14539>.
- [14] A. Varbella, B. Gjorgiev, G. Sansavini, "Geometric deep learning for online prediction of cascading failures in power grids," *Reliability Engineering & System Safety*, vol. 237, 2023, Art. no. 109341. <https://doi.org/10.1016/j.ress.2023.109341>.
- [15] F. T. Liu, K. M. Ting, and Z. Zhou, "Isolation forest," *IEEE International Conference on Data Mining*, 2008, pp. 413–422. <https://doi.org/10.1109/ICDM.2008.17>.
- [16] L. Busoniu *et al.*, "Reinforcement Learning for Control: Performance, Stability, and Deep Approximators," *Annual Reviews in Control*, vol. 46, pp. 8–28, 2018. <https://doi.org/10.1016/j.arcontrol.2018.09.005>.
- [17] S. Skogestad, "Control structure design for complete chemical plants," *Computers and Chemical Engineering*, vol. 28, no. 1–2, pp. 219–234, 2004. <https://doi.org/10.1016/j.compchemeng.2003.08.002>.
- [18] A. Kumar, R. Raj, A. Kumar, B. Verma, "Design of a novel mixed interval type-2 fuzzy logic controller for 2-DOF robot manipulator with payload," *Engineering Applications of Artificial Intelligence*, vol. 123, 2023, Art. no. 106329. <https://doi.org/10.1016/j.engappai.2023.106329>.
- [19] M. Raissi, P. Perdikaris, and G. Karniadakis, "Physics-informed neural networks: A deep learning framework for solving forward and inverse problems involving nonlinear partial differential equations," *Journal of Computational Physics*, vol. 378, pp. 686–707, 2019. <https://doi.org/10.1016/j.jcp.2018.10.045>.
- [20] A. Aziz, and Sowmyashree, "AI-Powered Predictive Maintenance System for Industrial Equipment: Reducing Downtime and Costs Through Sensor-B," *Journal of Scientific Research and Technology*, vol. 3, no. 8, 2025. <https://www.jsrtjournal.com/index.php/JSRT/article/view/387>.
- [21] A. Bozgeyik, L. Altay, A. Hepbasli, "Energetic, exergetic, exergoeconomic, environmental and sustainability analyses of a solar, geothermal and biomass based novel multi-generation system for production of power, hydrogen, heating, cooling and fresh water," *Process Safety and Environmental Protection*, vol. 177, 2023, pp. 400–415. <https://doi.org/10.1016/j.psep.2023.07.018>.
- [22] M. Shiea, A. Buffo, M. Vanni, D. L. Marchisio, "A novel finite-volume TVD scheme to overcome non-realizability problem in quadrature-based moment methods," *Journal of Computational Physics*, vol. 409, 2020, Art. no. 109337. <https://doi.org/10.1016/j.jcp.2020.109337>.
- [23] A. Morchid *et al.*, "Fire risk assessment system for food and sustainable farming using ai and IoT technologies: Benefits and challenges," *Internet of Things*, vol. 33, 2025, Art. no. 101704. <https://doi.org/10.1016/j.iot.2025.101704>.

- [24] E. G. Acuña Acuña, "Engineering-Oriented Bioinformatics Data Repository for Predictive Modeling of Complex Biological Systems". Zenodo, 2026. <https://doi.org/10.5281/zenodo.18190677>.
- [25] X. Meng, E. E. Karniadakis, "A composite neural network that learns from multi-fidelity data: Application to function approximation and inverse PDE problems," *Journal of Computational Physics*, vol. 401, 2020, Art. no. 109020. <https://doi.org/10.1016/j.jcp.2019.109020>.
- [26] I. Mezgár, J. Váncza, "From ethics to standards – A path via responsible AI to cyber-physical production systems," *Annual Reviews in Control*, vol. 53, 2022, pp. 391–404. <https://doi.org/10.1016/j.arcontrol.2022.04.002>.
- [27] S. Pan and Q. Yang, "A survey on transfer learning," *IEEE Transactions on Knowledge and Data Engineering*, vol. 22, no. 10, pp. 1345–1359, 2010. <https://doi.org/10.1109/TKDE.2009.191>.
- [28] S. M. Lundberg and S. I. Lee, "A unified approach to interpreting model predictions," *Advances in Neural Information Processing Systems*, 2017. <https://proceedings.neurips.cc/paper/2017/hash/8a20a8621978632d76c43dfd28b67767-Abstract.html>.
- [29] A. Haleem, M. Javaid, R. P. Singh, "Encouraging Safety 4.0 to enhance industrial culture: An extensive study of its technologies, roles, and challenges," *Green Technologies and Sustainability*, vol. 3, no. 3, 2025, Art. no. 100158. <https://doi.org/10.1016/j.grets.2024.100158>.
- [30] V. Salehi, H. Zarei, G. A. Shirali, K. Hajizadeh, "An entropy-based TOPSIS approach for analyzing and assessing crisis management systems in petrochemical industries," *Journal of Loss Prevention in the Process Industries*, vol. 67, 2020, Art. no. 104241. <https://doi.org/10.1016/j.jlp.2020.104241>.
- [31] D. Shan, B. Neis, "Employment-related mobility, regulatory weakness and potential fatigue-related safety concerns in short-sea seafaring on Canada's Great Lakes and St. Lawrence Seaway: Canadian seafarers' experiences," *Safety Science*, vol. 121, pp. 165-176, 2020. <https://doi.org/10.1016/j.ssci.2019.08.017>.
- [32] J. M. Douglas, *Conceptual Design of Chemical Processes*, New York, NY, USA: McGraw-Hill, 1988. <https://books.google.co.id/books?id=M6JTAAAAMAAJ>.
- [33] E. Glaessgen and D. Stargel, "The digital twin paradigm for future NASA and U.S. Air Force vehicles," in *Proc. 53rd AIAA Structures, Structural Dynamics and Materials Conf.*, 2012. <https://doi.org/10.2514/6.2012-1818>.
- [34] Q. Kong, K. Kuriyan, N. Shah, M. Guo, "Development of a responsive optimisation framework for decision-making in precision agriculture," *Computers & Chemical Engineering*, vol. 131, 2019, Art. no. 106585. <https://doi.org/10.1016/j.compchemeng.2019.106585>.
- [35] Z. Ge, Z. Song, S. X. Ding, and B. Huang, "Data mining and analytics in the process industry: The role of machine learning," *IEEE Access*, vol. 5, pp. 20590–20616, 2017. <https://doi.org/10.1109/ACCESS.2017.2756872>.
- [36] S. Yin, X. Zhu, and O. Kaynak, "Improved PLS focused on key-performance-indicator-related fault diagnosis," *IEEE Transactions on Industrial Electronics*, vol. 62, no. 3, pp. 1651–1658, 2015. <https://doi.org/10.1109/TIE.2014.2345331>.
- [37] J. An and S. Cho, "Variational autoencoder based anomaly detection using reconstruction probability," *Special Lecture on IE*, vol. 2, no. 1, pp. 1–18, 2015. <http://dm.snu.ac.kr/static/docs/TR/SNUDM-TR-2015-03.pdf>.
- [38] H. Yoo, B. Kim, J. W. Kim, J. H. Lee, "Reinforcement learning based optimal control of batch processes using Monte-Carlo deep deterministic policy gradient with phase segmentation," *Computers & Chemical Engineering*, vol. 144, 2021. <https://doi.org/10.1016/j.compchemeng.2020.107133>.
- [39] L. Yang, D. Zhang, and G. E. Karniadakis, "Physics-Informed Generative Adversarial Networks for Stochastic Differential Equations," *SIAM Journal on Scientific Computing*, vol. 42, no. 1, 2020. <https://doi.org/10.1137/18M1225409>.
- [40] Y. Xu *et al.*, "Physics-informed machine learning for reliability and systems safety applications: State of the art and challenges," *Reliability Engineering & System Safety*, vol. 230, 2023. <https://doi.org/10.1016/j.ress.2022.108900>.
- [41] M. Khani *et al.*, "Explainable AI reveals temporal risk pathways in fall prediction: Extracting clinical insights from multi-horizon machine learning models," *GeroScience*, 2026. <https://doi.org/10.1007/s11357-026-02117-x>.
- [42] E. G. Acuña Acuña, "AI-Ready Simulation Dataset for Chemical Process Safety and Accident Prevention". Zenodo, 2026. <https://doi.org/10.5281/zenodo.18200704>.

- [43] Z. A. B. Sinnah, “Explainable AI-driven predictive maintenance for mitigating process safety risks in safety-critical industrial equipment,” *J. Loss Prev. Process Ind.*, vol. 100, Art. no. 105907, Apr. 2026. <https://doi.org/10.1016/j.jlp.2025.105907>. <https://doi.org/10.1016/j.jlp.2025.105907>.
- [44] S. Mane *et al.*, “Digital twin in the chemical industry: A review,” *Digital Twins and Applications*, vol. 1, no. 2, pp. 118-130, 2024. <https://doi.org/10.1049/dgt2.12019>.
- [45] A. V. Alshehri, “Physics-Informed Neural Networks for Process Systems: Handling Plant-Model Mismatch,” *Industrial & Engineering Chemistry Research*, vol. 63, pp. 13650–13659, 2024. <https://doi.org/10.1021/acs.iecr.4c00690>.
- [46] H. Chen *et al.*, “Review of imbalanced fault diagnosis technology based on generative adversarial networks,” *Journal of Computational Design and Engineering*, vol. 11, no. 5, pp. 99–124, 2024. <https://doi.org/10.1093/jcde/qwae075>.
- [47] A. Hoenig, K. Roy, Y. T. Acquaah, S. Yi, and S. S. Desai, “Explainable AI for cyber-physical systems: Issues and challenges,” *IEEE Access*, vol. 12, pp. 73113–73140, 2024. <https://doi.org/10.1109/ACCESS.2024.3395444>.
- [48] B. Bolat-Akça and E. Bozkaya-Aras, “Digital twin-assisted intelligent anomaly detection system for Internet of Things,” *Ad Hoc Networks*, vol. 158, Art. no. 103484, May 2024. <https://doi.org/10.1016/j.adhoc.2024.103484>.