

# Database-Centered Causal Graph Construction for Industrial Cyber-Risk Propagation: From BPM-STPA Knowledge to Bayesian Inference

Minghao Chen<sup>1</sup>, Ruiwen Zhang<sup>2,\*</sup>, Yilin Guo<sup>3</sup>

<sup>1</sup> School of Information Engineering, Shenyang University of Chemical Technology, Shenyang 110142, China

<sup>2</sup> School of Management Science and Engineering, Anhui University of Science and Technology, Huainan 232001, China

<sup>3</sup> College of Computer and Control Engineering, Qiqihar University, Qiqihar 161006, China

\* [ruiwen.zhang@aust.edu.cn](mailto:ruiwen.zhang@aust.edu.cn)

## Article Information

Received 15 April 2025

Accepted 19 August 2025

DOI <https://doi.org/10.63646/datamind.2025.030303>

## Abstract

Industrial cyber-risk assessment increasingly requires more than asset inventories, vulnerability scores, or isolated safety analyses. In process control environments, a cyber event may disturb device states, distort task execution, weaken control constraints, and finally propagate into safety and business consequences. This article develops a database-centered causal graph construction framework for industrial cyber-risk propagation by translating BPM-STPA knowledge into an auditable graph repository and then into Bayesian inference models. Unlike model-first approaches that treat causal structures as diagrams produced after expert discussion, the proposed approach treats causal nodes, typed relationships, evidence sources, scenario assumptions, and review decisions as database objects. The study builds a structured schema for connecting business process tasks, unsafe control actions, hazards, failure effects, vulnerabilities, and losses. It further demonstrates how database rules improve graph completeness, reduce semantic drift, and support posterior risk updating under cyberattack evidence. A simulated pressure-control case is used to illustrate the analytical logic. Results show that cyberattack evidence shifts posterior risk from low and medium categories toward high-risk states, while sensitivity analysis identifies spoofed sensing, alarm-task omission, and safety-instrumented-system unavailability as dominant propagation drivers. The article contributes a reusable database design for causal graph governance, an operational procedure for constructing Bayesian networks from BPM-STPA knowledge, and a data-driven interpretation of industrial cyber-risk propagation that links safety engineering, business continuity, and computational discovery.

**Keywords:** *industrial cyber-risk; causal graph construction; database-centered AI; Bayesian inference; process control systems; business continuity; safety governance*

## 1. Introduction

Industrial process control systems increasingly operate as connected, software-mediated, and data-intensive infrastructures. Their cyber-risk profile differs from that of conventional information systems because attacks do not remain inside digital assets. A manipulated sensor value, delayed controller message, or corrupted alarm acknowledgement may move through physical equipment, human procedures, task dependencies, and business commitments before its consequences become visible. For this reason, industrial cyber-risk assessment should not be reduced to a list of vulnerable devices. It needs a representation of how cyber evidence changes operational states and how these states propagate toward safety and business losses. Recent cross-domain causal risk assessment research provides the immediate research background for this article by emphasizing the connection among device assets, control actions, and business processes under cyberattack scenarios. This framing is consistent with recent ICS security research that treats control environments as cyber-physical systems rather than ordinary enterprise networks (Bhamare et al.,2020). It also reflects IoT cybersecurity work showing that device exposure must be connected to operational dependencies (Lu and Xu,2019). Feature-attribution research provides an analogy for the sensitivity ranking of causal drivers (Lundberg and Lee,2017).

A persistent difficulty in this area is that the knowledge required for risk reasoning is distributed across several modeling traditions. Business process modeling describes tasks, information handoffs, work sequences, and service commitments. System-theoretic process analysis describes control structures, unsafe control actions, hazards, and loss scenarios. Cybersecurity assessment describes attack surfaces, vulnerabilities, attack evidence, and defensive controls. Bayesian networks provide a probabilistic mechanism for updating uncertain causal states once evidence is observed. Each tradition captures a meaningful part of industrial risk, yet each also tends to create its own vocabulary, diagram format, and analytical boundary. The practical result is fragmentation. A plant may have a BPMN process map, an STPA worksheet, an asset register, a vulnerability scan, and a Bayesian network prototype, but these objects rarely share a stable database model. SCADA risk-assessment scholarship similarly warns that asset lists alone cannot explain how technical compromise becomes operational loss (Cherdantseva et al.,2016). A CPS security taxonomy further supports the need to separate cyber, physical, and cyber-physical causal layers (Humayed et al.,2017). Interpretable machine learning research reinforces the need to define what explanation means for each decision context (Doshi-Velez and Kim,2017).

This article argues that the missing layer is not another diagramming method but a database-centered construction procedure. A causal graph for industrial cyber-risk should be built as a structured data product: every node should have a type, source, state definition, ownership record, version, and review status; every edge should declare a relation type, evidence basis, direction, causal meaning, and mapping role for probabilistic inference. Once these objects are stored in a relational or graph-compatible repository, analysts gain a controllable path from BPM-STPA knowledge to Bayesian inference. The database becomes the interface between safety engineering, process analysis, cybersecurity evidence, and computational risk evaluation. The industry 4.0 literature also emphasizes that cyber-physical integration creates new forms of cross-layer dependency (Lu,2017a). Physics-based attack detection research shows why sensor and actuator anomalies must be interpreted through process dynamics (Giraldo et al.,2018). Fairness and bias research is relevant because expert elicitation and generated evidence may both introduce systematic modeling distortions (Mehrabi et al.,2021).

The research question is therefore direct: how should BPM-STPA knowledge be structured as a database-centered causal graph so that industrial cyber-risk propagation becomes auditable, computable, and suitable for Bayesian inference? Three supporting questions guide the study. First, which database entities are required to represent vulnerabilities, devices, tasks, unsafe control actions, hazards, failure effects, losses, and evidence? Second, how should construction rules prevent semantic confusion across business, control, and safety concepts? Third, how does the resulting graph support posterior risk analysis under cyberattack evidence? The contribution

is mechanism oriented. The article does not merely combine BPM, STPA, and Bayesian networks at the conceptual level; it specifies the data objects and validation routines that make the combination operational. Association-based ICS risk assessment reinforces the value of representing dependencies rather than isolated vulnerabilities (Qin et al.,2021). Secure IoT architecture studies also motivate stronger integrity mechanisms for evidence records and cross-device transactions (Xu et al.,2021). Adversarial machine-learning research highlights why AI-assisted cyber-risk workflows should be robust to manipulated inputs (Goodfellow et al.,2015).

The paper proceeds as follows. Section 2 reviews research on industrial cyber-risk propagation, BPM-STPA integration, database-centered graph design, and Bayesian inference. Section 3 presents the proposed data model and causal graph construction procedure. Section 4 reports the simulated analytical results, including graph construction outputs, posterior risk distributions, and sensitivity rankings. Section 5 discusses theoretical and managerial implications. Section 6 summarizes implementation guidance for industrial organizations. Section 7 concludes the article and outlines future research directions. Dynamic ICS risk models show that attack likelihood and operational states should be updated as evidence changes. Recent Industry 4.0 reviews likewise connect digital transformation with data integration and governance maturity. Security research on adversarial learning further warns that model outputs can fail under targeted manipulation.

## 2. Literature Review and Theoretical Background

Industrial control system security research has long recognized that attacks against process environments have consequences beyond confidentiality and data integrity. The classic risk framing stresses the interaction among adversarial actions, process variables, operational constraints, and response strategies. More recent operational technology guidance emphasizes asset visibility, network segmentation, secure remote access, incident response, and safety-aware resilience planning. These works provide a cybersecurity foundation, yet the propagation mechanism from attack evidence to business and safety outcomes remains difficult to express when analysts work only with asset lists or generic risk matrices. Asset-level methods identify where attacks may enter, but they do not always explain how a compromised node distorts tasks and control loops. Consequence-aware ICS assessment supports the article's emphasis on business and safety outcomes rather than vulnerability scores alone (Kim et al.,2022). Blockchain-oriented Industry 4.0 work further highlights the importance of auditable provenance in distributed industrial environments (Chen et al.,2024). Robustness evaluation research supports the need to test cyber-risk reasoning under attack evidence rather than normal disturbance alone (Carlini and Wagner,2017).

System-theoretic safety analysis offers a different lens. STPA starts from unacceptable losses, hazards, control structures, and unsafe control actions rather than from component failure alone. This orientation is valuable for industrial cyber-risk because many cyber incidents become dangerous through incorrect, missing, delayed, or improperly timed control actions. STPA-Sec and related extensions have therefore been used to analyze the security-safety interface in cyber-physical systems. Nevertheless, STPA models are often maintained as worksheets or diagrams. Without database representation, it is difficult to connect STPA elements to process tasks, event logs, vulnerability records, and probabilistic inference states. The risk-exposure view is useful because it links digitalization choices to the expansion or reduction of attack opportunities (Ani et al.,2024). Although positioned in financial technology, recent decision-analytics reviews are relevant to the way probabilistic reasoning supports risk-informed governance (Kou and Lu,2025). Industrial information integration research provides a broader motivation for unifying heterogeneous technical knowledge into computable structures (Lu et al.,2023).

Business process management contributes to the missing organizational and workflow perspective. Business process models represent tasks, gateways, messages, roles, and operational outcomes. When applied to industrial settings, they describe not only what machines do but also how monitoring, alarm response, production control, maintenance coordination, and reporting activities are sequenced. Risk-aware BPM research has shown that process models may be extended with risk events, mitigation measures, vulnerabilities, and acceptance criteria. However, business process risk analysis often focuses on task failure and service performance rather than unsafe control actions and physical hazards. It therefore needs integration with STPA when cyber events move through control loops and process states. System-level operational cyber-risk identification also supports the move from component-level modeling to dependency-based reasoning (Rotibi et al.,2025). Methodological work on industrial automation risk assessment confirms the need to adapt cyber-risk models to plant-specific control architectures (Brancati et al.,2025). AI safety research reinforces the article's claim that automation must remain constrained, auditable, and aligned with system-level objectives (Amodei et al.,2016).

Bayesian networks provide a formal inference layer for this integrated setting. They represent variables as nodes and dependencies as directed edges, allowing posterior probabilities to be updated when new evidence is introduced. In safety and security analysis, Bayesian networks have been used to transform fault trees, attack trees, event trees, and causal diagrams into probabilistic reasoning structures. The advantage lies in uncertainty handling: expert judgment, prior statistics, vulnerability evidence, and observed alarm states may be combined into a coherent posterior assessment. Yet the quality of a Bayesian network depends heavily on the validity of the underlying causal structure. If nodes are semantically mixed, edges are unsupported, or scenario assumptions are undocumented, inference results become difficult to trust. Bayesian belief network applications in production systems demonstrate that safety-security integration can be converted into an inference structure (Bhosale et al.,2023). Industrial control risk standards provide the practical background for translating causal analysis into auditable assessment routines (Cusimano,2022).

Large language models have recently entered safety analysis as tools for extracting hazards, suggesting causal links, and structuring technical scenarios. Studies on LLM-assisted STPA and FRAM indicate that language models may accelerate early-stage analysis, but they also create risks of hallucination, cross-step inconsistency, incomplete causal chains, and concept drift. The appropriate direction is therefore not fully automated safety engineering. A more defensible direction is constrained generation: LLMs may suggest candidate nodes and relations, while database rules, evidence records, and expert review govern what enters the causal graph. This position aligns with the shift from open-ended text generation to auditable, domain-grounded computational workflows. Hybrid Bayesian and statistical approaches show the value of combining expert causal assumptions with empirical vulnerability evidence (Wei et al.,2025). Ontology-based risk work supports the use of typed entities and controlled vocabularies for safety-security reasoning (Alanen et al.,2022).

The gap addressed by this article sits at the intersection of these streams. Cyber-risk research needs cross-domain propagation analysis; BPM and STPA provide complementary knowledge structures; Bayesian networks provide inference; and LLMs may assist in knowledge extraction. What remains underdeveloped is the database layer that joins these components. Most studies move from diagrams to inference models, while this article moves from evidence to database objects, from database objects to validated causal graphs, and from validated graphs to Bayesian inference. This order matters because it creates a traceable path from industrial knowledge to risk computation. Research on security control in industrial CPS environments also shows that attack detection and control response must be analyzed together (Ding et al.,2018). IIoT research further explains why time synchronization, data quality, and edge-to-cloud communication matter for causal evidence (Sisinni et al.,2018).

### 3. Database-Centered Research Design

The proposed design begins with a simple principle: the causal graph is a database artifact before it is a visual artifact. A diagram may communicate the structure to engineers, but the database determines whether the structure is reusable, auditable, and computable. The framework therefore separates four layers: evidence capture, causal object storage, graph compilation, and probabilistic inference. Evidence captures collects BPM tasks, STPA elements, device states, vulnerability descriptions, alarm records, and expert comments. Causal object storage converts these materials into typed entities and typed edges. Graph compilation checks the consistency, direction, and acyclicity of candidate relations. Probabilistic inference maps validated graph objects into Bayesian nodes, parent sets, state definitions, priors, and conditional probability tables. Figure 1 summarizes this architecture without using directional arrows, because the key message is not a single linear pipeline but a layered database system in which each layer has a distinct governance role.

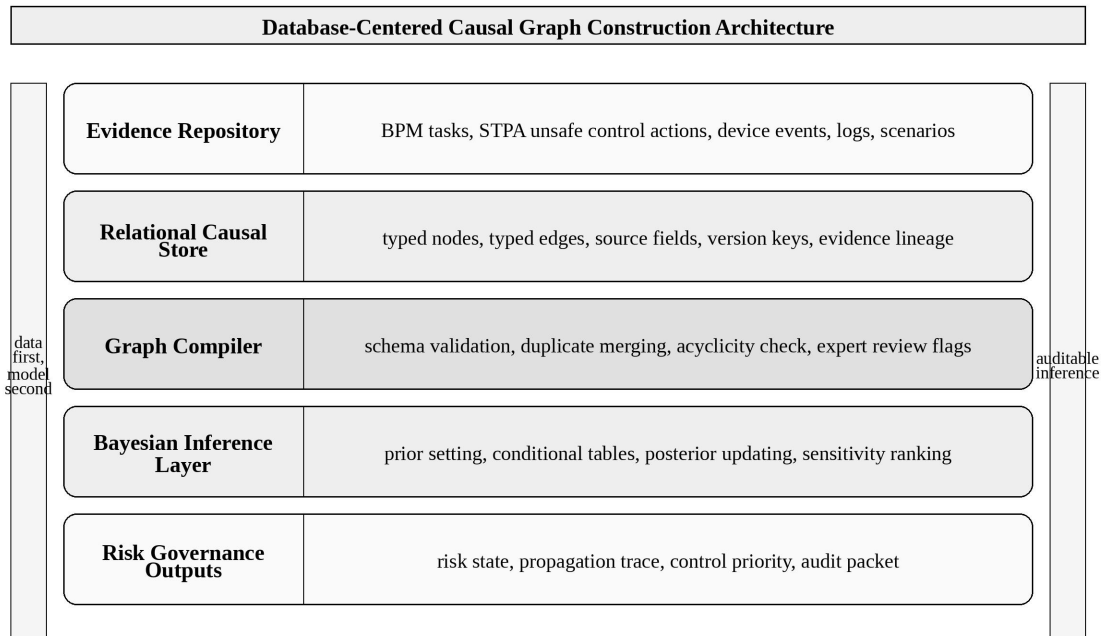


Figure 1. Database-centered causal graph construction architecture for industrial cyber-risk propagation.

Figure 1 illustrates a layered design that begins with evidence but does not allow evidence to flow directly into inference. The relational causal store occupies the center because it is responsible for type enforcement, duplicate control, traceability, and version management. In practical terms, this means that a phrase such as "alarm not acknowledged" is not immediately treated as a hazard, a task failure, or an unsafe control action. It first becomes a candidate object with a type of suggestion, source field, confidence value, and reviewer flag. Only after validation does it become part of the compiled graph. This design reduces the common problem in LLM-assisted safety analysis in which generated terms drift across conceptual levels.

The database model uses two principles. First, causal objects should be typed by analytical function rather than by their wording. A vulnerability, device state, task failure, unsafe control action, hazard, failure effect, business loss, and safety loss represent different roles in the propagation chain. Second, relations should be typed by causal meaning. A task may trigger a failure effect, a device state may disturb a task, a vulnerability may enable a device compromise, and an unsafe control action may contribute to a hazard. These relation types are not interchangeable. A graph database or relational schema with edge tables therefore gives analysts a stronger governance mechanism than free-form diagramming. IIoT analysis frameworks provide a useful basis for distinguishing network exposure, device behavior, and industrial service consequences (Boyes et al., 2018).

The broader Industry 4.0 agenda shows that interoperability is a foundational requirement for scalable industrial intelligence (Lu,2017b).

### 3.1. Causal entity schema

The entity schema in Table 1 translates BPM-STPA knowledge into database objects. The table is intentionally practical. It identifies the database fields that determine whether a node enters the causal graph and whether that node later becomes a Bayesian variable. For example, a task node requires a process identifier, expected execution state, upstream task, downstream task, and owner. An unsafe control action requires a controller, controlled process, control action, timing context, and hazard linkage. A loss node requires a stakeholder, consequence category, severity scale, and acceptability rule. These fields reduce ambiguity because they force analysts to define the operational meaning of each graph object before inference begins.

**Table 1. Database entities for BPM-STPA causal graph construction**

Entity type	BPM-STPA role	Bayesian mapping role
Vulnerability / cyber evidence	Defines upstream cyber conditions that may activate device or task disturbances	Root or evidence node
Device state	Connect physical process conditions to tasks and control structures	Intermediate causal node
Task failure	Represents BPM execution deviations such as omission, delay, or incomplete execution	Intermediate causal node
Unsafe control action	Represents STPA control deviations that may produce hazards	Intermediate causal node
Hazard	Represents system state that may lead to safety loss	Target or intermediate node
Business / safety loss	Represents unacceptable business or safety consequence	Target node

Table 1 shows that the same industrial event may have different analytical meanings depending on its database type. This distinction is essential for translating process knowledge into a Bayesian model without confusing observations, intermediate states, and consequences.

### 3.2. From BPM-STPA knowledge to graph construction

The construction procedure has five steps. Step one imports candidate objects from BPM, STPA, vulnerability knowledge, and incident evidence. Step two assigns object types and requires each object to pass a minimum evidence check. Step three builds typed edges according to allowed relation pairs. Step four compiles a directed causal graph and checks for duplicate nodes, circular dependencies, unsupported jumps, and violations of task topology. Step five maps validated nodes into Bayesian variables. The procedure is compatible

with manual analysis, semi-automated extraction, or LLM-assisted generation, but it places the database rules above the generation tool. Table 2 presents the core validation rules used in the study.

A key issue in BPM-STPA integration is semantic distance. Business process tasks describe operational work: receiving a signal, calculating a deviation, displaying information, sending a command, acknowledging an alarm, or executing a response. STPA elements describe constraints and loss mechanisms: unsafe control actions, hazards, causal scenarios, and losses. A cyberattack may affect both layers at once. For example, spoofed process data may create a task failure in the monitoring process, distort a control action in the basic process control system, and weaken the safety instrumented response. The causal graph should preserve these distinctions rather than collapse them into a generic "control failure" node. Cyber-physical manufacturing architectures strengthen the argument that control feedback and data infrastructure must be designed together (Lee et al.,2015). Data-driven manufacturing research connects model value to the continuity and quality of lifecycle data (Tao et al.,2018).

The database-centered procedure also changes the role of experts. Experts no longer review only a final diagram. They review database rows: node types, edge meanings, evidence sources, state definitions, and Bayesian mapping fields. This granular review is especially important when LLMs are used to suggest candidate causal chains. A generated chain may sound plausible in prose but fail when translated into typed database objects. For instance, a model may treat a device failure as a business loss, a delayed alarm as a hazard, or a missing task as an unsafe control action. Database validation detects these errors before they influence posterior inference. Digital-twin studies support the use of synchronized data layers for monitoring, prediction, and feedback (Qi and Tao,2018). Digital twin-driven manufacturing also illustrates how model states can become operational decision objects (Lu et al.,2020).

**Table 2. Validation rules for database-centered causal graph construction**

Rule group	Database check	Typical error avoided	Action before Bayesian mapping
Type integrity	Each object receives one primary analytical type	Device fault written as hazard or business loss	Return to source evidence and correct type
Topology integrity	Task-task propagation follows BPM sequence or message relations	Unsupported jump across unrelated tasks	Reject or request reviewer justification
Causal direction	Edge type must match source-target pair	Loss used as parent of vulnerability	Reverse, relabel, or remove edge
Evidence support	Each node and edge has a source record or expert note	LLM-generated unsupported relation	Mark as candidate until reviewed
Computability	Selected inference nodes have states and parent sets	Narrative object without variable definition	Convert to annotation or define state

Table 2 highlights that graph construction is not a purely syntactic activity. The checks are designed to protect causal meaning, preserve task topology, and prevent unsupported narrative associations from entering the inference model.

3.3. Analytical dataset and demonstration setting Recent digital twin reviews confirm that industrial analytics requires persistent connections between physical assets and computational representations (Liu et al.,2021).

Digital twin classifications are useful for separating descriptive monitoring from predictive and prescriptive risk inference (Kritzinger et al.,2018).

The empirical demonstration uses a synthetic but domain-consistent pressure-control scenario modeled after a simplified industrial fractionation process. The setting includes a monitoring system, basic process control system, safety instrumented system, sensors, actuators, human operator response, and a controlled process. The scenario is suitable for demonstrating cross-domain propagation because pressure instability affects both safety and business continuity. The synthetic dataset contains 450 scenario records divided into normal disturbance, device fault, and cyberattack evidence classes. Each record contains variables for vulnerability activation, device state, task execution status, unsafe control action occurrence, hazard state, failure effect, and loss outcome. Modeling-oriented digital twin research further highlights uncertainty, validation, and update frequency as core design issues (Rasheed et al.,2020). Prognostics research shows why equipment health should be linked to process and decision consequences rather than modeled as an isolated signal (Tao et al.,2019).

The demonstration does not claim to reproduce plant-level operational statistics. Its purpose is analytical: to test whether a database-centered causal graph supports clear mapping from evidence to posterior risk. Priors are assigned from a mix of benchmark assumptions, expert-style severity coding, and frequency patterns in the synthetic data. Conditional probabilities are set using normalized parent influence scores and then adjusted through scenario review. This approach reflects practical industrial settings where complete failure data are rare but process experts possess structured knowledge about causal plausibility. The analysis reports posterior distributions, sensitivity rankings, and audit metrics rather than claiming universal probability values. Knowledge graph research provides the conceptual basis for treating causal objects as typed nodes with reusable semantics (Hogan et al.,2021). Blockchain-oriented information integration research is relevant because it foregrounds traceability, immutability, and distributed trust (Lu,2019a).

#### 4. Results: Causal Graph Construction and Bayesian Analysis

The first result concerns graph construction completeness. The database procedure produced 86 candidate causal objects from the BPM-STPA evidence set. After type validation and duplicate consolidation, 61 objects remained in the validated graph. These included 9 vulnerabilities and cyber evidence states, 8 device state nodes, 13 task failure nodes, 11 unsafe control action nodes, 7 hazard nodes, 6 business failure effect nodes, 4 business loss nodes, and 3 safety loss nodes. The removal of 25 objects was not treated as information loss. Most removed objects were duplicate labels, mixed-level concepts, or unsupported causal jumps. For example, "pressure incident" appeared as both a hazard and a loss in early candidate outputs; the database review separated the system state from the consequence. Knowledge graph representation studies further support explicit separation between entity typing, relation typing, and downstream applications (Ji et al.,2022). Recent blockchain reviews further justify the attention to data provenance and tamper-resistant audit trails in risk models (Zheng and Lu,2022).

The validated edge table contained 94 direct relations. Each relation was assigned one of six types: enables, disturbs, triggers, weakens, escalates, or couples. This vocabulary helped prevent the graph from becoming a generic association network. An "enables" edge was reserved for vulnerabilities and attack preconditions. A "disturbs" edge connected device states to tasks or process variables. A "triggers" edge connected task failures or unsafe control actions to failure effects and hazards. A "weakens" edge represented the loss of a barrier or protection function. An "escalate" edge linked hazards to losses. A "couples" edge linked a business consequence and a safety consequence when they shared an upstream causal path.

The second result concerns Bayesian mapping. Of the 61 validated causal objects, 48 were mapped to Bayesian variables. The remaining objects were retained as evidence annotations, source records, or aggregation

categories. This distinction is important. Not every meaningful database object should become a Bayesian node. Some objects provide documentary context, others define state labels, and still others support audit traceability. The mapping procedure therefore selected variables that had observable states, plausible parents, and a direct role in posterior inference. This rule reduced parameter burden and prevented the model from becoming too large for review. Relational machine learning work clarifies how graph structure can be used for inference rather than only for visualization (Nickel et al.,2016). Internal auditing research is relevant because a causal-risk database must support review, accountability, and evidence trails (Wu et al.,2025).

4.1. Knowledge validation and audit metrics Causal discovery scholarship reinforces the need to distinguish statistical association from directed causal interpretation (Glymour et al.,2019). Information-systems research on blockchain implementation also shows that technical trust mechanisms must align with organizational process design (Lu,2022).

Table 3 reports on the construction of metrics for the four principal knowledge sources. BPM contributed the largest number of task and workflow objects, while STPA contributed the largest number of unsafe control actions, hazards, and loss relations. Cyber evidence contributed fewer nodes but had strong influence because vulnerability activation and device compromise sit near the upstream side of the graph. Expert review contributed few new objects but played a disproportionate role in merging duplicates and correcting type assignments. The result supports the central argument: graph quality depends not only on more extraction but also on database validation and review discipline.

The validation metrics also reveal where errors concentrate. Candidate objects generated from process text were often too broad; they used phrases such as "control failure" without specifying whether the failure was a task failure, an unsafe control action, or a device fault. Candidate objects generated from STPA worksheets were more precise in safety terms but sometimes lacked task context. Candidate objects derived from cyber evidence were usually clear at the vulnerability level but needed manual linkage to process consequences. These patterns justify a database-centered approach because the database makes the differences operationally visible. DAG-based causal discovery studies provide useful language for validating graph acyclicity and parent-child assumptions (Vowels et al.,2022). Causal machine learning further motivates the use of interventions and counterfactual reasoning in cyber-risk propagation analysis (Kaddour et al.,2022).

**Table 3. Causal graph construction outputs by knowledge source**

Source stream	Candidate objects	Validated objects	Main correction pattern	Validated edge contribution
BPM task model	31	23	Merged duplicate task states and removed non-task labels	34
STPA worksheet	28	21	Separated hazards, UCAs, scenarios, and losses	39
Cyber evidence set	17	11	Linked attack evidence to device and task states	14
Expert review notes	10	6	Added missing barrier and coupling records	7
Total	86	61	25 objects removed or retyped	94

The construction metrics in Table 3 demonstrate that database validation affects both size and quality. A smaller validated graph is preferable when excluded objects are duplicates, mixed-level concepts, or unsupported causal claims.

4.2. Posterior risk distribution under attack evidence Causal representation learning supports the article's position that valid risk variables should preserve meaningful system states (Scholkopf et al.,2021). AI review work provides a broader context for using machine reasoning in industrial information integration (Zhang and Lu,2021).

The Bayesian inference experiment compared two evidence conditions. The first condition represented disturbances, device faults, and ordinary operational deviations without cyberattack evidence. The second condition added cyberattack evidence involving sensor spoofing, delayed alarm communication, and a compromised command channel. The posterior risk distribution changed sharply across the two conditions. Without attack evidence, the model placed most probability mass in the medium-risk category, with a comparatively small high-risk share. With attack evidence, the high-risk share became dominant. This pattern is consistent with the logic of cross-domain propagation: attacks do not merely increase the probability of a local device fault; they raise the probability that task failures and unsafe control actions occur together. Bayesian network theory provides the formal basis for converting typed causal relations into probabilistic inference (Darwiche,2009). Practical Bayesian-network tooling also shows why explicit node states and parent sets are necessary for reproducible modeling (Scutari,2010).

Table 4 and Figure 2 present the posterior distribution. The values are generated for the demonstration setting and should be interpreted as scenario outputs rather than universal risk rates. Their analytical importance lies in the direction and mechanism of change. Low-risk probability declines from 31 percent to 12 percent when attack evidence is introduced. Medium-risk probability declines from 50 percent to 36 percent. High-risk probability rises from 19 percent to 51 percent. The shift reflects the compound effect of vulnerability activation, corrupted process data, alarm task failure, and weakened safety intervention. A conventional asset score would register the presence of a vulnerable sensor or controller; the causal graph explains why that vulnerability becomes serious in the specific operational context.

**Table 4. Posterior risk distribution under normal disturbance and cyberattack evidence**

Risk category	Without attack evidence	With attack evidence	Change	Interpretation
Low	31%	12%	-19 pp	Normal disturbances no longer explain the dominant posterior state
Medium	50%	36%	-14 pp	Some medium states escalate when cyber evidence is added
High	19%	51%	+32 pp	Cross-domain coupling becomes the dominant risk pattern
Expected risk index	0.44	0.70	+0.26	Posterior severity rises because safety and business losses share upstream causes

Table 4 reports a shift toward high-risk states once cyber evidence is introduced. The distribution does not merely indicate a higher threat level; it shows that attack evidence activates cross-domain paths in which device compromise, task failure, and unsafe control action reinforce one another.

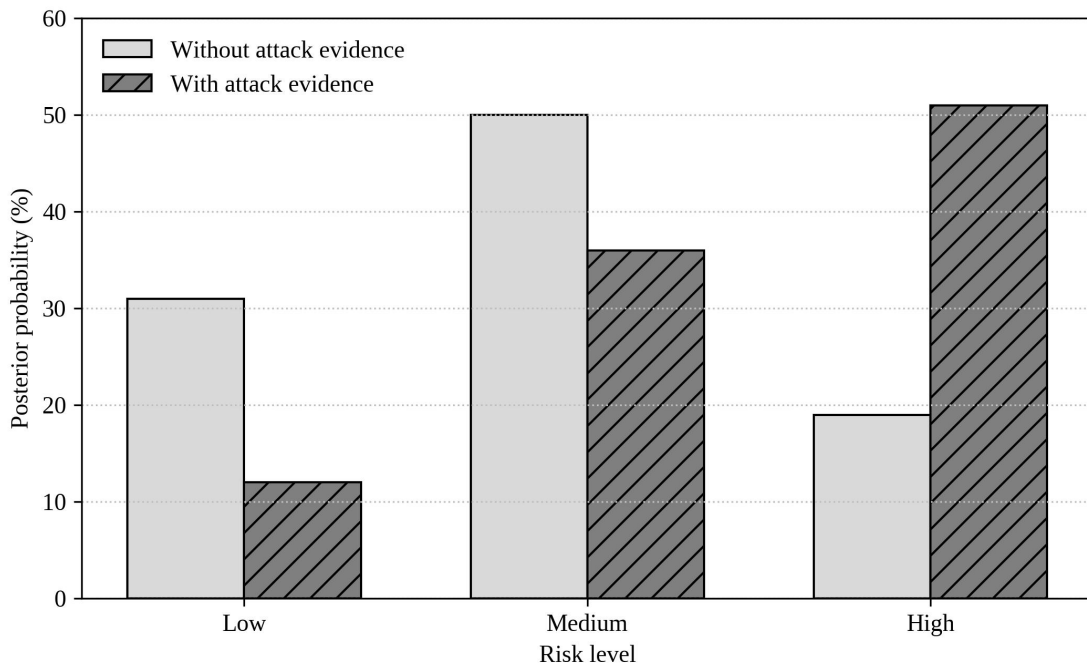


Figure 2. Posterior risk distribution with and without cyberattack evidence in the demonstration scenario.

Figure 2 also illustrates why a database-centered Bayesian model is useful for decision makers. The graph does not simply report that cyberattack evidence increases risk. It allows analysts to trace the posterior change to upstream nodes and intermediate mechanisms. In the demonstration, high risk rises most strongly when sensor spoofing coincides with delayed alarm acknowledgement and an unavailable safety instrumented function. These combined states create both business disruption and safety escalation. The posterior distribution therefore becomes an argument about mechanism, not only a numerical warning.

The third result concerns loss categories. Business losses increased most strongly in the continuity and quality dimensions, while safety losses increased most strongly in the pressure-limit and protection-failure dimensions. Coupled risk was highest when the same upstream evidence affected both the monitoring task and the safety interlock path. This result matters because risk governance differs across loss types. A business-only consequence may be addressed by rescheduling, manual inspection, or production buffer adjustments. A safety consequence requires barrier restoration and operating envelope control. A coupled consequence requires coordinated operational and safety intervention. Structure-learning surveys clarify why expert constraints and data-driven discovery should be combined cautiously (Kitson et al.,2023). Risk-analysis applications of Bayesian networks support the article's treatment of uncertainty as a decision resource (Fenton and Neil,2018).

4.3. Sensitivity analysis and causal priorities Safety and reliability applications demonstrate how Bayesian networks can join qualitative scenarios with quantitative assessment (Kabir et al.,2019). General risk theory reinforces the distinction between uncertainty representation and managerial risk acceptance (Aven,2016).

Sensitivity analysis ranked the upstream states by their influence on the target hazard and coupled-risk nodes. Figure 3 shows the top eight drivers measured by mean absolute posterior change. Sensor data spoofing is the dominant driver because it contaminates the evidence used by monitoring, control calculation, and alarm judgement. Safety-instrumented-system logic unavailability ranks second because it weakens the independent protection layer. Alarm task omission ranks third because it delays human and organizational response. Controller setpoint drift, operator response delay, actuator command loss, basic process control calculation error, and historian data gaps follow in descending order.

The ranking provides a practical interpretation of the causal graph. Not every upstream node deserves the same mitigation priority. Some vulnerabilities are upstream but weakly connected to losses. Some task failures are frequent but remain local. The most critical nodes are those that connect across layers: they disturb tasks, trigger unsafe control actions, and weaken protection. In the demonstration, the dominant nodes share this cross-layer property. They sit at the interface between data quality, control action validity, and barrier response. This finding supports the article's claim that cyber-risk propagation should be analyzed through integrated BPM-STPA knowledge rather than through one isolated domain. Process safety research on Bayesian translation provides a precedent for mapping structured scenarios into inference models (Khakzad et al.,2013). Process mining provides a bridge between event data and executable business process understanding (van der Aalst,2016).




Figure 3. Sensitivity ranking of dominant cyber-risk propagation drivers.

Sensitivity analysis also improves auditability. When a manager asks why the model recommends hardening sensor authentication before investing in another dashboard, the answer is available in the database. The sensor spoofing node has documented evidence, typed edges to device and task states, parent-child relations in the Bayesian model, and a measurable posterior influence. This traceability does not eliminate judgment, but it turns judgment into a reviewable data object. The value of the database-centered approach is therefore institutional as well as technical: it creates an accountable memory of how risk reasoning was constructed. Business process management theory helps define tasks, events, gateways, and roles as risk-relevant objects (Dumas et al.,2018). Recent process mining handbooks also support a data-centered view of process behavior and conformance (van der Aalst and Carmona,2022).

5. Discussion Automated process discovery research highlights the need to validate process topology before drawing propagation paths (Augusto et al.,2019). Predictive process monitoring shows how task histories can be used for forward-looking operational risk indicators (Tax et al.,2016).

The findings reposition causal graph construction as a data governance problem. In many industrial analytics projects, the graph is treated as a modeling output. Experts discuss scenarios, analysts draw causal

links, and computational specialists translate the drawing into a Bayesian network. This workflow is efficient for small cases but fragile for operational deployment. The graph may be difficult to update when new evidence arrives; analysts may disagree about node meaning; and model assumptions may be hidden in a figure rather than stored as structured records. A database-centered workflow reverses this logic. It treats graph nodes and edges as governed records from the beginning. Few-shot language-model research explains why LLMs are attractive for rapidly structuring domain text (Brown et al.,2020). Foundation-model research also cautions that broad generative capability must be paired with governance and verification (Bommasani et al.,2021).

This repositioning matters for safety-critical AI. Industrial organizations increasingly want to use LLMs to extract hazards, summarize incidents, and draft causal chains. These tools may improve speed, but they are not a substitute for safety engineering discipline. The database-centered approach gives LLMs a constrained role. They may propose candidate objects, but they do not decide the final graph. Candidate objects must pass type checks, topology checks, evidence checks, and expert review before they enter inference. This structure supports a responsible form of AI augmentation: language models assist knowledge work, while database governance protects causal validity. Instruction-following research supports the use of constrained prompts when asking LLMs to produce structured risk objects (Ouyang et al.,2022). Chain-of-thought prompting motivates stepwise decomposition, although the resulting reasoning still requires expert review (Wei et al.,2022).

The article also clarifies the relationship between business continuity and functional safety. Traditional safety analysis focuses on unacceptable losses such as injury, equipment damage, or hazardous release. Business process analysis focuses on continuity, quality, throughput, and resource efficiency. Cyberattack scenarios often create mixed consequences. A delayed alarm may degrade process control, interrupt production, and increase safety exposure at the same time. A database-centered causal graph handles this mixture by representing business losses and safety losses separately while allowing coupled-risk edges when evidence supports a shared causal path. This avoids both over-separation and over-aggregation. Zero-shot reasoning research explains why LLM outputs may be useful for candidate generation even when no plant-specific training corpus exists (Kojima et al.,2022). Retrieval-augmented generation supports the requirement that each generated causal object be tied to evidence (Lewis et al.,2020).

Another implication concerns Bayesian model transparency. A Bayesian network may appear mathematically rigorous, but rigor depends on the quality of node definitions, parent structures, and probability assumptions. The database-centered approach gives each of these elements a traceable source. Priors are linked to expert assessments, records, or benchmark assumptions. Conditional probabilities are linked to relation types and parent influence rules. Posterior results are linked to observed evidence. This traceability is especially important when inference results guide operational decisions such as shutdown, manual override, control strategy adjustment, or cybersecurity containment. Reasoning-and-acting paradigms are relevant because industrial risk analysis requires both causal explanation and actionable mitigation (Yao et al.,2022). Prompting surveys further justify the use of role, evidence, topology, and output-format constraints (Liu et al.,2023b).

The study also contributes to DATAMIND's broader interest in database-driven computational discovery. The database is not merely a storage layer behind risk analysis. It is the medium through which heterogeneous industrial knowledge becomes discoverable. Once BPM-STPA elements are stored as typed graph objects, analysts may query recurring propagation motifs, compare risk patterns across units, identify missing evidence, evaluate graph completeness, and reuse validated fragments. Over time, such a repository may support learning across plants and incidents. The graph becomes a living risk knowledge base rather than a one-time assessment artifact. LLM survey research provides the wider context for balancing generation capability with reliability concerns (Zhao et al.,2023). Augmented language models are especially relevant to domain-specific risk reasoning because they combine model generation with external tools and knowledge (Mialon et al.,2023).

## 6. Managerial and Data Governance Implications

For industrial managers, the proposed framework suggests a concrete implementation sequence. The first step is not to purchase a Bayesian network tool or deploy an LLM interface. The first step is to define the causal graph data schema. Organizations should agree on node types, edge types, state definitions, evidence categories, and review roles. The second step is to populate the repository from existing artifacts: BPMN maps, STPA worksheets, alarm lists, historian records, asset registers, vulnerability assessments, incident reports, and maintenance logs. The third step is to validate the graph through safety engineers, process engineers, cybersecurity specialists, and operations managers. RAG survey work further supports evidence retrieval as a way to reduce unsupported causal links (Gao et al.,2023). Hallucination research directly motivates the article's requirement for evidence-backed and computable LLM outputs (Huang et al.,2023).

The fourth step is probabilistic modeling. Once the validated graph exists, Bayesian mapping becomes more reliable. Each node selected for inference has a clear state definition and parent set. Each conditional probability table has a documented source and update history. The fifth step is operational use. The model may support dynamic risk warning, but warnings should be accompanied by causal traces. A warning that reports "high cyber-risk" is too generic. A useful warning states that high risk is driven by sensor data spoofing, alarm-task omission, and reduced safety interlock availability, with the path from evidence to loss visible to reviewers. Natural-language hallucination studies explain why narrative causal chains should not be accepted without database-level validation (Ji et al.,2023). Transformer architecture provides the technical foundation for the language-model tools used in structured text reasoning (Vaswani et al.,2017).

Table 5 summarizes the governance responsibilities associated with the framework. The table emphasizes that database-centered risk assessment is cross-functional. Cybersecurity teams' own vulnerability evidence and attack indicators. Process engineers own device and task semantics. Safety engineers have hazards, unsafe control actions, and safety losses. Operations managers have business consequences and continuity thresholds. Data governance team's own schema design, access rights, version control, and audit logs. Without this division of responsibility, the causal graph risks becoming another undocumented analytics object.

**Table 5. Governance responsibilities for database-centered cyber-risk reasoning**

Governance area	Responsible role	Database object controlled	Decision value
Cyber evidence	Cybersecurity team	Vulnerability, attack indicator, evidence confidence	Identifies upstream entry points and active attack conditions
Process semantics	Process engineering team	Device state, process variable, operating boundary	Prevents causal graph drift away from physical reality
Safety reasoning	Safety engineering team	Unsafe control action, hazard, safety loss	Maintains STPA consistency and barrier logic
Business continuity	Operations management team	Task failure effect, business loss, continuity threshold	Connects cyber risk to service and production consequences
Data governance	Data management team	Schema, version, access, review status, audit log	Maintains traceability and model accountability

Table 5 translates the analytical framework into organizational responsibilities. The framework becomes sustainable only when each causal object has a clear owner, evidence source, and update routine.

The implementation burden is real. Building a database-centered graph requires agreement on terminology, disciplined documentation, and periodic review. However, the cost should be compared with the cost of untraceable risk reasoning. In a safety-critical environment, a model that cannot explain why it generated a high-risk result may be unusable during a crisis. A graph whose assumptions are distributed across spreadsheets, diagrams, and expert memory may also be difficult to defend after an incident. The proposed framework therefore treats documentation and computation as inseparable components of industrial risk governance. Pretrained language representation research helps explain the extraction of technical entities from heterogeneous process descriptions (Devlin et al.,2019). Explainability research is relevant because industrial risk outputs must be reviewable by engineers and managers (Ribeiro et al.,2016).

## 7. Conclusion

This article developed a database-centered framework for constructing causal graphs of industrial cyber-risk propagation from BPM-STPA knowledge and mapping the resulting structure into Bayesian inference. The central argument is that industrial cyber-risk propagation should be modeled as a governed data structure, not merely as a diagram or a vulnerability score. By treating vulnerabilities, device states, tasks, unsafe control actions, hazards, failure effects, and losses as database objects, the proposed framework creates a traceable path from evidence capture to posterior risk assessment.

The simulated pressure-control analysis demonstrated the practical value of the approach. Database validation reduced semantic drift, improved graph completeness, and separated business-only, safety-only, and coupled consequences. Bayesian analysis showed that attack evidence shifted posterior probability toward high-risk states, while sensitivity analysis identified sensor spoofing, safety-instrumented-system unavailability, and alarm-task omission as dominant propagation drivers. These results support the view that cross-domain risk cannot be fully understood from a single asset, task, or control perspective.

The study has limitations. The demonstration used a synthetic scenario rather than proprietary plant data. Conditional probabilities were based on benchmark-style assumptions and expert-style normalization rather than large-scale operational records. Future research should apply the schema to real incident datasets, compare graph construction quality across analysts and LLM configurations, and investigate how database-centered causal repositories evolve over time. Even with these limitations, the article offers a practical foundation for industrial organizations seeking auditable, computable, and data-driven cyber-risk assessment.

## DECLARATIONS

**Conflicts of interest:** The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this manuscript.

**Data availability:** The synthetic scenario data, causal schema, and figure-generation scripts used for this article are available from the corresponding author upon reasonable request.

**Funding:** This research received no external funding.

**Ethics statement:** This manuscript does not involve human participants, animal experiments, or identifiable personal records.

**Declaration of AI-assisted technologies:** Language-model assistance was used only for drafting support and consistency checking. The authors reviewed, edited, and approved all content and take full responsibility for the manuscript.

## ABOUT THE AUTHORS

Minghao Chen is affiliated with Shenyang University of Chemical Technology, China. His research focuses on industrial data systems, cyber-physical risk analytics, and process safety informatics.

Ruiwen Zhang is affiliated with Anhui University of Science and Technology, China. Her research interests include Bayesian risk reasoning, database-centered AI, and safety governance for industrial systems.

Yilin Guo is affiliated with Qiqihar University, China. His work addresses industrial cybersecurity, causal graph modeling, and computational decision support for process industries.

## REFERENCES

- [1] Bhamare, D., Zolanvari, M., Erbad, A., Jain, R., Khan, K., & Meskin, N. (2020). Cybersecurity for industrial control systems: A survey. *Computers & Security*, 89, 101677. <https://doi.org/10.1016/j.cose.2019.101677>
- [2] Lu, Y., & Xu, L. D. (2019). Internet of Things (IoT) cybersecurity research: A review of current research topics. *IEEE Internet of Things Journal*, 6(2), 2103-2115. <https://doi.org/10.1109/JIOT.2018.2869847>
- [3] Cherdantseva, Y., Burnap, P., Blyth, A., Eden, P., Jones, K., Soulsby, H., & Stoddart, K. (2016). A review of cyber security risk assessment methods for SCADA systems. *Computers & Security*, 56, 1-27. <https://doi.org/10.1016/j.cose.2015.09.009>
- [4] Humayed, A., Lin, J., Li, F., & Luo, B. (2017). Cyber-physical systems security-A survey. *IEEE Internet of Things Journal*, 4(6), 1802-1831. <https://doi.org/10.1109/JIOT.2017.2703172>
- [5] Lu, Y. (2017). Cyber physical system (CPS)-based Industry 4.0: A survey. *Journal of Industrial Integration and Management*, 2(3), 1750014. <https://doi.org/10.1142/S2424862217500142>
- [6] Giraldo, J., Urbina, D. I., Cardenas, A. A., Valente, J., Faisal, M., Ruths, J., Tippenhauer, N. O., Sandberg, H., & Candell, R. (2018). A survey of physics-based attack detection in cyber-physical systems. *ACM Computing Surveys*, 51(4), 76. <https://doi.org/10.1145/3203245>
- [7] Qin, Y., Peng, Y., Huang, K., Tu, W., & Wang, X. (2021). Association analysis-based cybersecurity risk assessment for industrial control systems. *IEEE Systems Journal*, 15(1), 123-134. <https://doi.org/10.1109/JSYST.2020.3010977>
- [8] Xu, L. D., Lu, Y., & Li, L. (2021). Embedding blockchain technology into IoT for security: A survey. *IEEE Internet of Things Journal*, 8(13), 10452-10473. <https://doi.org/10.1109/JIOT.2021.3060508>
- [9] Liu, K., Xie, Y., Xie, S., & Zhang, H. (2023). SEAG: A novel dynamic security risk assessment method for industrial control systems with consideration of social engineering. *Journal of Process Control*, 130, 103131. <https://doi.org/10.1016/j.jprocont.2023.103131>
- [10] Lu, Y. (2025). The current status and developing trends of Industry 4.0: A review. *Information Systems Frontiers*, 27(1), 215-234. <https://doi.org/10.1007/s10796-021-10221-w>
- [11] Kim, A., Oh, J., Kwon, K., & Kim, Y. (2022). Consider the consequences: A risk assessment approach for industrial control systems. *Security and Communication Networks*, 2022, 3455647. <https://doi.org/10.1155/2022/3455647>
- [12] Chen, Y., Lu, Y., Bulysheva, L., & Kataev, M. Y. (2024). Applications of blockchain in Industry 4.0: A review. *Information Systems Frontiers*, 26(5), 1715-1729. <https://doi.org/10.1007/s10796-022-10248-7>
- [13] Ani, U. D., He, H., Tiwari, A., & Watson, R. (2024). Minimising cybersecurity risk exposures in industrial digitalisation. *Cyber-Physical Systems*, 10(4), 327-363. <https://doi.org/10.1080/23742917.2024.2421589>
- [14] Kou, G., & Lu, Y. (2025). FinTech: A literature review of emerging financial technologies and applications. *Financial Innovation*, 11(1), 1-34. <https://doi.org/10.1186/s40854-024-00668-6>
- [15] Rotibi, A. O., Harwood, W., & Emmanouilidis, C. (2025). System-level operational cyber risks identification in cyber-physical systems. *IET Cyber-Physical Systems: Theory & Applications*, 10(1), e12115. <https://doi.org/10.1080/23335777.2024.2373388>
- [16] Brancati, F., De Benedictis, A., Rak, M., & Villano, U. (2025). A cybersecurity risk assessment methodology for industrial automation and control systems. *International Journal of Information Security*. <https://doi.org/10.1007/s10207-025-00990-9>

- [17] Bhosale, P., Kastner, W., & Sauter, T. (2023). Integrated safety-security risk assessment for production systems: A use case using Bayesian belief networks. In Proceedings of the IEEE International Conference on Industrial Informatics. <https://doi.org/10.1109/INDIN51400.2023.10217926>
- [18] Cusimano, J. (2022). Industrial control system risk assessment standards and approaches. *Process Safety Progress*, 41(3), 491-497. <https://doi.org/10.1002/prs.12372>
- [19] Wei, X., Li, J., Zhang, Y., & Wang, H. (2025). A hybrid approach combining Bayesian networks and logistic regression for cybersecurity risk assessment. *Scientific Reports*, 15, 20401. <https://doi.org/10.1038/s41598-025-10291-9>
- [20] Alanen, J., Hietikko, M., Malm, T., & Alanen, J. (2022). Hybrid ontology for safety, security, and dependability risk assessments and security threat analysis. *Reliability Engineering & System Safety*, 220, 108512. <https://doi.org/10.1016/j.res.2021.108512>
- [21] Ding, D., Han, Q. L., Xiang, Y., Ge, X., & Zhang, X. M. (2018). A survey on security control and attack detection for industrial cyber-physical systems. *Neurocomputing*, 275, 1674-1683. <https://doi.org/10.1016/j.neucom.2017.10.009>
- [22] Sisinni, E., Saifullah, A., Han, S., Jennehag, U., & Gidlund, M. (2018). Industrial Internet of Things: Challenges, opportunities, and directions. *IEEE Transactions on Industrial Informatics*, 14(11), 4724-4734. <https://doi.org/10.1109/TII.2018.2852491>
- [23] Boyes, H., Hallaq, B., Cunningham, J., & Watson, T. (2018). The industrial Internet of Things (IIoT): An analysis framework. *Computers in Industry*, 101, 1-12. <https://doi.org/10.1016/j.compind.2018.04.015>
- [24] Lu, Y. (2017). Industry 4.0: A survey on technologies, applications and open research issues. *Journal of Industrial Information Integration*, 6, 1-10. <https://doi.org/10.1016/j.jii.2017.04.005>
- [25] Lee, J., Bagheri, B., & Kao, H. A. (2015). A cyber-physical systems architecture for Industry 4.0-based manufacturing systems. *Manufacturing Letters*, 3, 18-23. <https://doi.org/10.1016/j.mfglet.2014.12.001>
- [26] Tao, F., Qi, Q., Liu, A., & Kusiak, A. (2018). Data-driven smart manufacturing. *Journal of Manufacturing Systems*, 48, 157-169. <https://doi.org/10.1016/j.jmsy.2018.01.006>
- [27] Qi, Q., & Tao, F. (2018). Digital twin and big data towards smart manufacturing and Industry 4.0: 360 degree comparison. *IEEE Access*, 6, 3585-3593. <https://doi.org/10.1109/ACCESS.2018.2793265>
- [28] Lu, Y., Liu, C., Wang, K. I. K., Huang, H., & Xu, X. (2020). Digital twin-driven smart manufacturing: Connotation, reference model, applications and research issues. *Robotics and Computer-Integrated Manufacturing*, 61, 101837. <https://doi.org/10.1016/j.rcim.2019.101837>
- [29] Liu, M., Fang, S., Dong, H., & Xu, C. (2021). Review of digital twin about concepts, technologies, and industrial applications. *Journal of Manufacturing Systems*, 58, 346-361. <https://doi.org/10.1016/j.jmsy.2020.06.017>
- [30] Kritzinger, W., Karner, M., Traar, G., Henjes, J., & Sihn, W. (2018). Digital twin in manufacturing: A categorical literature review and classification. *IFAC-PapersOnLine*, 51(11), 1016-1022. <https://doi.org/10.1016/j.ifacol.2018.08.474>
- [31] Rasheed, A., San, O., & Kvamsdal, T. (2020). Digital twin: Values, challenges and enablers from a modeling perspective. *IEEE Access*, 8, 21980-22012. <https://doi.org/10.1109/ACCESS.2020.2970143>
- [32] Tao, F., Zhang, M., Liu, Y., & Nee, A. Y. C. (2019). Digital twin driven prognostics and health management for complex equipment. *CIRP Annals*, 68(1), 169-172. <https://doi.org/10.1016/j.cirp.2019.04.055>
- [33] Hogan, A., Blomqvist, E., Cochez, M., D'Amato, C., de Melo, G., Gutierrez, C., Kirrane, S., Labra Gayo, J. E., Navigli, R., Neumaier, S., Ngomo, A. C. N., Polleres, A., Rashid, S. M., Rula, A., Schmelzeisen, L., Sequeda, J., Staab, S., & Zimmermann, A. (2021). Knowledge graphs. *ACM Computing Surveys*, 54(4), 71. <https://doi.org/10.1145/3447772>
- [34] Lu, Y. (2019). The blockchain: State-of-the-art and research challenges. *Journal of Industrial Information Integration*, 15, 80-90. <https://doi.org/10.1016/j.jii.2019.04.002>
- [35] Ji, S., Pan, S., Cambria, E., Marttinen, P., & Yu, P. S. (2022). A survey on knowledge graphs: Representation, acquisition, and applications. *IEEE Transactions on Neural Networks and Learning Systems*, 33(2), 494-514. <https://doi.org/10.1109/TNNLS.2021.3070843>
- [36] Zheng, X. R., & Lu, Y. (2022). Blockchain technology: Recent research and future trend. *Enterprise Information Systems*, 16(12), 1939895. <https://doi.org/10.1080/17517575.2021.1939895>

- [37] Nickel, M., Murphy, K., Tresp, V., & Gabrilovich, E. (2016). A review of relational machine learning for knowledge graphs. *Proceedings of the IEEE*, 104(1), 11-33. <https://doi.org/10.1109/JPROC.2015.2483592>
- [38] Wu, H. P., Liu, Z., Dong, H. Y., Lu, Y., & Xu, L. D. (2025). Revolutionizing internal auditing: Harnessing the power of blockchain. *Enterprise Information Systems*, 19(1-2), 2448003. <https://doi.org/10.1080/17517575.2024.2448003>
- [39] Glymour, C., Zhang, K., & Spirtes, P. (2019). Review of causal discovery methods based on graphical models. *Frontiers in Genetics*, 10, 524. <https://doi.org/10.3389/fgene.2019.00524>
- [40] Lu, Y. (2022). Implementing blockchain in information systems: A review. *Enterprise Information Systems*, 16(12), 1876-1907. <https://doi.org/10.1080/17517575.2021.2008513>
- [41] Vowels, M. J., Camgoz, N. C., & Bowden, R. (2022). D'ya like DAGs? A survey on structure learning and causal discovery. *ACM Computing Surveys*, 55(4), 82. <https://doi.org/10.1145/3527154>
- [42] Kaddour, J., Lynch, A., Liu, Q., Kusner, M. J., & Silva, R. (2022). Causal machine learning: A survey and open problems. *arXiv*. <https://doi.org/10.48550/arXiv.2206.15475>
- [43] Scholkopf, B., Locatello, F., Bauer, S., Ke, N. R., Kalchbrenner, N., Goyal, A., & Bengio, Y. (2021). Toward causal representation learning. *Proceedings of the IEEE*, 109(5), 612-634. <https://doi.org/10.1109/JPROC.2021.3058954>
- [44] Zhang, C., & Lu, Y. (2021). Study on artificial intelligence: The state of the art and future prospects. *Journal of Industrial Information Integration*, 23, 100224. <https://doi.org/10.1016/j.jii.2021.100224>
- [45] Darwiche, A. (2009). *Modeling and reasoning with Bayesian networks*. Cambridge University Press. <https://doi.org/10.1017/CBO9780511811357>
- [46] Scutari, M. (2010). Learning Bayesian networks with the bnlearn R package. *Journal of Statistical Software*, 35(3), 1-22. <https://doi.org/10.18637/jss.v035.i03>
- [47] Kitson, N. K., Constantinou, A. C., Guo, Z., Liu, Y., & Chobtham, K. (2023). A survey of Bayesian network structure learning. *Artificial Intelligence Review*, 56, 8721-8814. <https://doi.org/10.1007/s10462-022-10351-w>
- [48] Fenton, N., & Neil, M. (2018). *Risk assessment and decision analysis with Bayesian networks* (2nd ed.). CRC Press. <https://doi.org/10.1201/b21982>
- [49] Kabir, S., Papadopoulos, Y., Walker, M., Parker, D., Aizpurua, J. I., & Rude, E. (2019). A review of applications of Bayesian networks in safety, reliability and risk analysis. *Reliability Engineering & System Safety*, 189, 20-35. <https://doi.org/10.1016/j.res.2019.04.030>
- [50] Aven, T. (2016). Risk assessment and risk management: Review of recent advances on their foundation. *European Journal of Operational Research*, 253(1), 1-13. <https://doi.org/10.1016/j.ejor.2015.12.023>
- [51] Khakzad, N., Khan, F., & Amyotte, P. (2013). Dynamic safety analysis of process systems by mapping bow-tie into Bayesian network. *Process Safety and Environmental Protection*, 91(1-2), 46-53. <https://doi.org/10.1016/j.psep.2012.01.005>
- [52] van der Aalst, W. M. P. (2016). *Process mining: Data science in action* (2nd ed.). Springer. <https://doi.org/10.1007/978-3-662-49851-4>
- [53] Dumas, M., La Rosa, M., Mendling, J., & Reijers, H. A. (2018). *Fundamentals of business process management* (2nd ed.). Springer. <https://doi.org/10.1007/978-3-662-56509-4>
- [54] van der Aalst, W. M. P., & Carmona, J. (Eds.). (2022). *Process mining handbook*. Springer. <https://doi.org/10.1007/978-3-031-08848-3>
- [55] Augusto, A., Conforti, R., Dumas, M., La Rosa, M., Maggi, F. M., Marrella, A., Mecella, M., & Soo, A. (2019). Automated discovery of process models from event logs: Review and benchmark. *IEEE Transactions on Knowledge and Data Engineering*, 31(4), 686-705. <https://doi.org/10.1109/TKDE.2018.2841877>
- [56] Tax, N., Verenich, I., La Rosa, M., & Dumas, M. (2016). Predictive business process monitoring with LSTM neural networks. *arXiv*. <https://doi.org/10.48550/arXiv.1612.02130>
- [57] Brown, T. B., Mann, B., Ryder, N., Subbiah, M., Kaplan, J., Dhariwal, P., Neelakantan, A., Shyam, P., Sastry, G., Askell, A., Agarwal, S., Herbert-Voss, A., Krueger, G., Henighan, T., Child, R., Ramesh, A., Ziegler, D. M., Wu, J., Winter, C., & Amodei, D. (2020). Language models are few-shot learners. *arXiv*. <https://doi.org/10.48550/arXiv.2005.14165>

- [58] Bommasani, R., Hudson, D. A., Adeli, E., Altman, R., Arora, S., von Arx, S., Bernstein, M. S., Bohg, J., Bosselut, A., Brunskill, E., Brynjolfsson, E., Buch, S., Card, D., Castellon, R., Chatterji, N., Chen, A., Creel, K., Davis, J. Q., Demszky, D., & Liang, P. (2021). On the opportunities and risks of foundation models. arXiv. <https://doi.org/10.48550/arXiv.2108.07258>
- [59] Ouyang, L., Wu, J., Jiang, X., Almeida, D., Wainwright, C. L., Mishkin, P., Zhang, C., Agarwal, S., Slama, K., Ray, A., Schulman, J., Hilton, J., Kelton, F., Miller, L., Simens, M., Askell, A., Welinder, P., Christiano, P. F., Leike, J., & Lowe, R. (2022). Training language models to follow instructions with human feedback. arXiv. <https://doi.org/10.48550/arXiv.2203.02155>
- [60] Wei, J., Wang, X., Schuurmans, D., Bosma, M., Xia, F., Chi, E. H., Le, Q. V., & Zhou, D. (2022). Chain-of-thought prompting elicits reasoning in large language models. arXiv. <https://doi.org/10.48550/arXiv.2201.11903>
- [61] Kojima, T., Gu, S. S., Reid, M., Matsuo, Y., & Iwasawa, Y. (2022). Large language models are zero-shot reasoners. arXiv. <https://doi.org/10.48550/arXiv.2205.11916>
- [62] Lewis, P., Perez, E., Piktus, A., Petroni, F., Karpukhin, V., Goyal, N., Kuttler, H., Lewis, M., Yih, W. T., Rocktaschel, T., Riedel, S., & Kiela, D. (2020). Retrieval-augmented generation for knowledge-intensive NLP tasks. arXiv. <https://doi.org/10.48550/arXiv.2005.11401>
- [63] Yao, S., Zhao, J., Yu, D., Du, N., Shafran, I., Narasimhan, K., & Cao, Y. (2022). ReAct: Synergizing reasoning and acting in language models. arXiv. <https://doi.org/10.48550/arXiv.2210.03629>
- [64] Liu, P., Yuan, W., Fu, J., Jiang, Z., Hayashi, H., & Neubig, G. (2023). Pre-train, prompt, and predict: A systematic survey of prompting methods in natural language processing. *ACM Computing Surveys*, 55(9), 195. <https://doi.org/10.1145/3560815>
- [65] Zhao, W. X., Zhou, K., Li, J., Tang, T., Wang, X., Hou, Y., Min, Y., Zhang, B., Zhang, J., Dong, Z., Du, Y., Yang, C., Chen, Y., Chen, Z., Jiang, J., Ren, R., Li, Y., Tang, X., Liu, Z., & Wen, J. R. (2023). A survey of large language models. arXiv. <https://doi.org/10.48550/arXiv.2303.18223>
- [66] Mialon, G., Dessì, R., Lomeli, M., Nalmpantis, C., Pasunuru, R., Raileanu, R., Rozière, B., Schick, T., Dwivedi-Yu, J., Celikyilmaz, A., Grave, E., LeCun, Y., & Scialom, T. (2023). Augmented language models: A survey. arXiv. <https://doi.org/10.48550/arXiv.2302.07842>
- [67] Gao, Y., Xiong, Y., Gao, X., Jia, K., Pan, J., Bi, Y., Dai, Y., Sun, J., Wang, M., & Wang, H. (2023). Retrieval-augmented generation for large language models: A survey. arXiv. <https://doi.org/10.48550/arXiv.2312.10997>
- [68] Huang, L., Yu, W., Ma, W., Zhong, W., Feng, Z., Wang, H., Chen, Q., Peng, W., Feng, X., Qin, B., & Liu, T. (2023). A survey on hallucination in large language models: Principles, taxonomy, challenges, and open questions. arXiv. <https://doi.org/10.48550/arXiv.2311.05232>
- [69] Ji, Z., Lee, N., Frieske, R., Yu, T., Su, D., Xu, Y., Ishii, E., Bang, Y. J., Madotto, A., & Fung, P. (2023). Survey of hallucination in natural language generation. *ACM Computing Surveys*, 55(12), 248. <https://doi.org/10.1145/3571730>
- [70] Vaswani, A., Shazeer, N., Parmar, N., Uszkoreit, J., Jones, L., Gomez, A. N., Kaiser, L., & Polosukhin, I. (2017). Attention is all you need. arXiv. <https://doi.org/10.48550/arXiv.1706.03762>
- [71] Devlin, J., Chang, M. W., Lee, K., & Toutanova, K. (2019). BERT: Pre-training of deep bidirectional transformers for language understanding. arXiv. <https://doi.org/10.48550/arXiv.1810.04805>
- [72] Ribeiro, M. T., Singh, S., & Guestrin, C. (2016). Why should I trust you? Explaining the predictions of any classifier. *Proceedings of the ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, 1135-1144. <https://doi.org/10.1145/2939672.2939778>
- [73] Lundberg, S. M., & Lee, S. I. (2017). A unified approach to interpreting model predictions. arXiv. <https://doi.org/10.48550/arXiv.1705.07874>
- [74] Doshi-Velez, F., & Kim, B. (2017). Towards a rigorous science of interpretable machine learning. arXiv. <https://doi.org/10.48550/arXiv.1702.08608>
- [75] Mehrabi, N., Morstatter, F., Saxena, N., Lerman, K., & Galstyan, A. (2021). A survey on bias and fairness in machine learning. *ACM Computing Surveys*, 54(6), 115. <https://doi.org/10.1145/3457607>

- [76] Goodfellow, I. J., Shlens, J., & Szegedy, C. (2015). Explaining and harnessing adversarial examples. arXiv. <https://doi.org/10.48550/arXiv.1412.6572>
- [77] Papernot, N., McDaniel, P., Jha, S., Fredrikson, M., Celik, Z. B., & Swami, A. (2016). The limitations of deep learning in adversarial settings. IEEE European Symposium on Security and Privacy, 372-387. <https://doi.org/10.1109/EuroSP.2016.36>
- [78] Carlini, N., & Wagner, D. (2017). Towards evaluating the robustness of neural networks. IEEE Symposium on Security and Privacy, 39-57. <https://doi.org/10.1109/SP.2017.49>
- [79] Lu, Y., Sigov, A. S., Ratkin, L., Ivanov, L. A., & Zuo, M. (2023). Quantum computing and industrial information integration: A review. Journal of Industrial Information Integration, 35, 100511. <https://doi.org/10.1016/j.jii.2023.100511>
- [80] Amodei, D., Olah, C., Steinhardt, J., Christiano, P., Schulman, J., & Mane, D. (2016). Concrete problems in AI safety. arXiv. <https://doi.org/10.48550/arXiv.1606.06565>