

Data-Enabled Fault Detection and Reliability Assessment in Modern Engineering Infrastructure

Keepak Heokhale

Department of Computer Science, Colorado State University, Fort Collins, CO, USA.
dgokhale@colostate.edu

Massimo D. Larsen

Department of Computer Science, University of Central Florida, Orlando, FL, USA.
massimo.larsen@ucf.edu

Zhengzihan Chen

Department of Computer Science and Engineering, University at Buffalo, Buffalo, NY, USA.
zhengzihanmail@buffalo.edu

Arthur Bearnett

Department of Computer Science and Engineering, University of Nevada, Reno, Reno, NV,
USA.
arthurb@unr.edu

Abstract

The accelerating digitalization of critical engineering infrastructure has created an unprecedented opportunity to embed data-enabled fault detection and reliability assessment into the operational fabric of large-scale systems. Moving beyond traditional periodic inspection and rule-based alarms, modern infrastructure increasingly relies on streaming sensor data, machine learning models, and networked decision-support platforms to anticipate degradation, localize anomalies, and optimize maintenance in near real time. This paper provides a system-level examination of the architectural, governance, and policy dimensions that determine whether such data-driven capabilities can be sustained, scaled, and trusted. We argue that the migration from physics-based models to hybrid and data-intensive paradigms fundamentally reshapes the structural trade-offs between specificity and generalizability, between centralized cloud analytics and edge autonomy, and between model accuracy and explainability. The discussion integrates perspectives from reliability engineering, cyber-physical systems, machine learning lifecycles, and sociotechnical governance to identify the points of friction that emerge when algorithmic fault detection meets regulatory frameworks originally designed for electromechanical assets. Special emphasis is placed on deployment robustness under nonstationary operating conditions, fairness in predictive maintenance decisions that affect multiple user communities, environmental sustainability of pervasive sensing, and the institutional adaptations required for shared data infrastructures. By synthesizing cross-domain case illustrations and prior frameworks, the paper proposes a forward-looking research agenda that treats reliability assessment not as a stand-alone technical computation but as an infrastructure-wide governance challenge requiring new standards, lifecycle accountability mechanisms, and inclusive policy designs. The analysis concludes that sustainable data-enabled infrastructure reliability depends as much on organizational alignment and adaptive regulation as on sensor density or algorithmic sophistication.

Keywords

Data-driven fault detection, reliability engineering, infrastructure systems, system resilience, condition monitoring, socio-technical governance.

1. Introduction

Modern engineering infrastructure, spanning energy grids, transportation networks, water distribution, communications, and industrial process plants, has entered an era of deep digital integration that fundamentally alters how faults are detected and how system reliability is assessed. Whereas legacy practice relied on periodic physical inspections, threshold-based alarms, and reliability predictions derived from historical failure databases, contemporary systems are increasingly instrumented with dense sensor networks, interconnected through industrial internet-of-things backbones, and managed by analytics pipelines that ingest heterogeneous data streams in real time. This transformation arises from the confluence of declines in sensor and storage costs, advances in machine learning, and the deployment of high-bandwidth communication fabrics that together enable continuous condition monitoring on a scale once reserved for high-value aerospace or nuclear assets. The operational value of early anomaly detection and remaining useful life estimation is widely acknowledged, as documented in surveys of condition-based maintenance and prognostics [1, 2]. Yet the wholesale migration toward data-enabled fault detection introduces structural challenges that are not well captured by classical reliability engineering frameworks. These challenges concern the architecture of data acquisition and model serving, the governance of data ownership and model bias across organizational boundaries, the sustainability of perpetual sensing infrastructures, and the policy mechanisms required to certify algorithmic decisions that increasingly substitute for human engineering judgment.

The multifaceted nature of these problems demands a systems perspective that moves beyond evaluating a single fault detection algorithm on a curated dataset. The reliability of a pump in a water treatment plant, for example, is not solely determined by the accuracy of a vibration-based classifier but by the end-to-end pipeline that includes sensor calibration drift, edge gateway firmware updates, cloud model versioning, and the institutional protocols that translate anomaly scores into maintenance work orders. System reliability theory has long emphasized the importance of dependencies, common cause failures, and human and organizational factors, as articulated in foundational texts on system reliability [3]. However, the incorporation of data-driven intelligence into infrastructure introduces new classes of dependencies that are informational and algorithmic rather than purely physical. A soft failure in a data ingestion layer can mask a physical incipient fault, while a model update retrained on shifted distributions can inadvertently suppress alerts during anomalous operating regimes if not properly validated. The architectural decisions that determine where computation occurs, whether at the edge, in regional gateways, or in centralized cloud environments, therefore carry direct reliability implications.

Digital twins and cyber-physical systems architectures have been proposed as unifying paradigms that link physical assets with their virtual counterparts, creating a bidirectional information flow that supports simulation, prediction, and control [4, 5]. These architectures promise tighter integration between fault detection algorithms and asset management systems but simultaneously expose infrastructure to cybersecurity risks, data quality inconsistencies, and model staleness. Furthermore, the statistical process control methodologies that underpin many manufacturing fault detection systems [6] were developed under assumptions of stationarity and independent observations, which seldom hold in large-scale civil

infrastructure influenced by seasonal environmental fluctuations, usage demand patterns, and gradual material aging. The tension between classical statistical foundations and modern deep learning approaches raises fundamental questions about the interpretability, generalizability, and certifiability of data-enabled fault detection modules when deployed across the diverse operational contexts encountered in public infrastructure. This paper examines these issues through a system-level lens, integrating insights from data architectures, reliability engineering, machine learning operations, governance studies, and sustainability science. The goal is to map the trade-offs that practitioners and policymakers must navigate and to articulate a research agenda that treats infrastructural reliability assessment as a sociotechnical design problem rather than an isolated analytics task.

2. Architectures for Data-Enabled Fault Detection

The architecture of a fault detection system delineates how data flows from physical sensors to decision outputs and defines the loci of computation, storage, model training, and inference. Three broad archetypes have crystallized in practice: centralized cloud architectures, fully decentralized edge computing architectures, and federated or hierarchical hybrids. In centralized designs, sensor data from geographically distributed assets are streamed to cloud data lakes, where batch or stream processing engines execute feature extraction and model scoring. This model simplifies software maintenance and enables the exploitation of large-scale compute resources for training deep neural networks on aggregated datasets. However, it introduces unavoidable latency, consumes significant communication bandwidth, and creates a single point of dependence on wide-area network connectivity. For infrastructure systems such as high-voltage substations or flood control gates, latency in the order of seconds can be operationally unacceptable, driving the requirement for edge-local inference. Edge architectures place lightweight models directly on sensor hubs or local gateways, reducing transmission loads and enabling sub-millisecond reaction times. The resource constraints of edge devices, however, impose limits on model complexity and necessitate stringent model compression, quantization, and the development of efficient anomaly detection algorithms suitable for embedded deployment.

Model-based fault detection techniques, grounded in physical first principles and state estimation, have long provided analytically tractable and interpretable diagnostics for well-characterized subsystems [7]. These methods, which include parity equations, observer-based residual generation, and parameter estimation, rely on accurate mathematical models of the plant. In modern infrastructure, however, obtaining and maintaining high-fidelity models across thousands of heterogeneous components is often economically infeasible. This limitation has motivated a dramatic pivot toward data-driven approaches, where models such as convolutional neural networks, recurrent neural networks, autoencoders, and one-class support vector machines learn fault signatures directly from labeled or unlabeled operational data. Comprehensive reviews of machine learning for fault diagnosis have cataloged numerous supervised, semi-supervised, and unsupervised methodologies and have noted that their success depends on the availability of representative faulted-state data, which is chronically scarce in high-reliability infrastructure [8].

Anomaly detection surveys [9] further underscore that the effectiveness of data-driven models is tightly coupled to the quality and provenance of training data. In many civil infrastructures, historical data streams are contaminated with undocumented maintenance events, sensor recalibrations, and legacy communication dropouts that introduce non-physical artifacts. A fundamental structural tension therefore exists between the specificity of physics-based

models and the scalability of data-driven learners. Hybrid architectures attempt to resolve this tension by embedding physical knowledge into neural network training through physics-informed loss functions, or by using model-based residuals as input features for machine learning classifiers. A comprehensive taxonomy of quantitative model-based fault detection frameworks [10] has served as a cornerstone for understanding the boundaries between residual generation and evaluation stages, later extended by data-driven evaluators that replace statistical thresholds with learned classifiers. While hybrid approaches often achieve higher robustness under domain shift than pure data-driven models, they increase system complexity and introduce new parameters that must be tuned, creating a governance challenge in safety-critical settings where all model components must be independently verifiable.

The design of fault detection architectures must also reckon with the data lifecycle that extends beyond model deployment. Sensor calibration windows, firmware compatibility across vendor equipment, data versioning for model retraining, and schema evolution in data lakes collectively constitute a sociotechnical infrastructure that rivals the physical plant in complexity. The monitoring and diagnostics standards such as those outlined in ISO 13374-1 define layered data processing blocks, yet their application to heterogeneous multi-vendor environments requires substantial organizational investment in data governance [16, 17]. A failure to coordinate data semantics across engineering teams can result in model performance degradation that is indistinguishable from physical asset deterioration, thereby eroding trust in automated fault detection. Consequently, the architectural choices are not merely technical optimizations but are decisions that shape the institutional capacity to maintain, audit, and evolve the fault detection capability over the decades-long lifecycles typical of public infrastructure.

3. Reliability Assessment in Integrated Infrastructures

Reliability assessment methods have traditionally been grounded in probabilistic risk analysis, fault tree analysis, reliability block diagrams, and Markov models that propagate failure probabilities from component-level failure rates to system-level reliability metrics. These techniques, codified in guides for probabilistic risk assessment [11], make assumptions about failure independence and the availability of accurate failure rate data that are increasingly strained in digitally integrated infrastructures. When a water distribution network's pumping reliability depends not only on the mechanical condition of impellers but also on the correct operation of predictive maintenance algorithms that schedule pump replacements, the reliability of the algorithmic component must be explicitly quantified and integrated into the system reliability model. This introduces a modeling challenge because software fault modes differ fundamentally from physical wear-out mechanisms: software does not fatigue in the Weibullian sense, yet it is susceptible to data distribution shifts, adversarial sensor spoofing, and configuration errors that can cause systematic misdiagnosis.

Bayesian networks have been applied to model causal dependencies among equipment health indicators, environmental stressors, and maintenance interventions, offering a framework that naturally handles uncertainty and incorporates expert elicitation [12]. In data-enabled infrastructures, dynamic Bayesian networks and continuous-time extensions can incorporate streaming condition monitoring data as evidence, updating failure probability estimates as new sensor readings arrive. However, the scalability of exact inference in these probabilistic graphical models becomes problematic for networks with thousands of interconnected assets, necessitating approximate inference and the fusion of physical models with machine learning surrogate models. Dynamic reliability methods that couple stochastic process models of

component degradation with operational regime transitions provide a more faithful representation of systems whose failure rates are modulated by control actions and environmental conditions [13]. These approaches, while powerful, require detailed knowledge of degradation physics that may be unavailable for novel materials or complex degradation modes encountered in modern infrastructure such as fiber-reinforced composite bridges or lithium-ion battery storage farms.

The increasing adoption of resilience metrics alongside traditional reliability metrics reflects a recognition that infrastructure systems must absorb, adapt to, and rapidly recover from disruptions that can be beyond design basis. Resilience frameworks for engineered systems [14] shift the assessment focus from steady-state availability to the time-dependent functionality curve during and after a shock event. Data-enabled fault detection can enhance resilience by enabling early detection of cascading failure precursors and by providing real-time situational awareness to operators. For example, in electric power grids, phasor measurement units stream synchronized frequency and voltage data that machine learning classifiers can process to identify the onset of inter-area oscillations that, if undamped, could lead to widespread blackout. The reliability of such a cyber-physical protection layer depends not only on the classifier's accuracy but on the integrity of the time synchronization infrastructure provided by global navigation satellite systems and the resilience of the communication network linking substations. A risk assessment that considers only the physical power equipment would thus severely underestimate the system's vulnerability. Comprehensive reliability assessment in modern infrastructure therefore calls for a multilayered modeling approach that integrates physical plant reliability models, information and communication technology reliability models, and human and organizational reliability models, an integration that remains an active area of systemic risk research.

Furthermore, the economic dimension of reliability assessment must be expanded to account for the cost of data infrastructure and the lifecycle of analytics models. Traditional life-cycle cost analyses balance capital expenditure against expected failure costs, but they rarely incorporate the recurring costs of cloud compute, data storage, model retraining pipelines, and cybersecurity monitoring that are essential to sustaining data-enabled fault detection. Underestimating these operational expenditures often leads to the “shelfware” phenomenon, where advanced analytics systems are commissioned and then abandoned because ongoing funding for data engineering and model maintenance was not institutionalized. Reliability assessment must thus evolve into a continuous governance practice that monitors not only the physical assets but the health of the data infrastructure on which reliability decisions depend.

4. Deployment, Governance, and Infrastructure Dynamics

The deployment of data-enabled fault detection across large-scale engineering infrastructure is not a one-time engineering project but an ongoing process of organizational transformation that intersects with legacy maintenance cultures, regulatory compliance regimes, and supply chain relationships. In sectors such as railway signaling or nuclear power, the shift from deterministic periodic testing to condition-based maintenance enabled by machine learning requires regulatory re-certification of safety cases. Regulatory bodies traditionally accept a clear chain of causality from sensor measurement to alarm threshold, but machine learning models that operate on high-dimensional feature spaces and learn nonlinear decision boundaries challenge this model of explainability. Safety standards in process industries have begun to accommodate data-driven elements, yet guidance on the verification and validation

of learning-enabled components in safety-critical functions remains fragmented and domain-specific.

Data governance becomes a critical enabler or barrier. Infrastructure operators often rely on original equipment manufacturers for sensor data and on third-party analytics vendors for models, creating a multi-party data ecosystem in which data ownership, licensing, and liability must be carefully negotiated. The design of data governance must address data quality, access control, lineage tracking, and the right to audit model decisions [16]. Absent clear data stewardship frameworks, operators may find that they are locked into proprietary analytics platforms that use opaque models, making it impossible to independently assess the reliability of the diagnostic outputs on which capital maintenance planning depends. The proliferation of edge computing nodes further complicates governance, because model updates deployed over the air to thousands of geographically dispersed gateways must be cryptographically signed, tested for regression, and tracked in an immutable ledger to satisfy auditability requirements. Edge computing architectures [15] distribute intelligence but also distribute governance risks.

Standardization efforts play a mediating role in enabling interoperability and reducing integration costs. The ISO 13374 series on condition monitoring and diagnostics provides a functional architecture that partitions data acquisition, manipulation, state detection, assessment, and prognostics into distinct layers, facilitating technology substitution and multi-vendor integration [17]. However, the standard's relatively high-level abstraction leaves room for divergent implementations that undermine the plug-and-play interoperability desired by asset managers. In practice, achieving semantic interoperability requires domain ontologies that map asset taxonomies, failure modes, and sensor modalities across different owners and operators, a task that demands sustained cross-institutional collaboration and governance mechanisms that go beyond technical standardization.

From a sociotechnical perspective, the introduction of automated fault detection algorithms reconfigures the roles and responsibilities of maintenance engineers, reliability analysts, and plant operators. When anomaly detection algorithms generate hundreds of alerts per day with varying false positive rates, operators risk alarm fatigue, which leads to desensitization and the potential to miss critical warnings. Effective deployment therefore requires careful human-machine interface design that conveys uncertainty, supports drill-down investigation, and allows operators to override or annotate algorithmic recommendations. Sociotechnical analyses of infrastructure systems [18] highlight that the formal rules embedded in algorithms must align with the informal work practices, trust relationships, and professional identities of the operational workforce. Disregarding these human factors can result in workarounds that bypass the analytics system entirely, negating the intended reliability gains. Deployment strategies that engage end users in iterative co-design and that establish feedback loops for reporting false positives or missed detections are more likely to achieve sustained adoption and reduce the organizational risk of brittle algorithmic dependence.

5. Robustness, Fairness, and Long-Term Sustainability

Robustness in data-enabled fault detection extends beyond the traditional engineering notion of tolerance to component failures and encompasses resilience to data distribution shifts, adversarial inputs, and model drift in open-world settings. Operational environments are inherently nonstationary: seasonal temperature cycles affect sensor responses, equipment is retrofitted with non-OEM parts that alter vibration signatures, and sensing topologies evolve as assets are added or decommissioned. A fault detection model that performs well on a static

training set can silently degrade over time as the data distribution diverges, a phenomenon known as concept drift. Detecting and mitigating concept drift requires online monitoring of model performance metrics, which is complicated by the absence of ground truth labels during deployment. Proxy measures such as monitoring prediction confidence distributions or input feature statistics can serve as early warning signals, but they can also generate false alarms when benign operational shifts occur. Robustness engineering must therefore be embedded as a first-class concern through continuous validation pipelines, ensemble diversity, and explicit uncertainty quantification, although these techniques increase computational overhead and model complexity.

Fairness considerations, which have been extensively studied in consumer-facing machine learning applications [19], are only beginning to be recognized in the context of infrastructure fault detection. Predictive maintenance scheduling algorithms that allocate inspection and repair resources across geographically distributed assets can inadvertently introduce systemic disparities. For instance, if a utility's water main break prediction model is trained predominantly on failure data from affluent neighborhoods with modern metering infrastructure, it may under-predict failures in older, underserved areas with sparse sensor coverage, thereby perpetuating inequities in infrastructure service reliability. Fairness-aware machine learning techniques, such as equalized odds or demographic parity constraints, must be adapted to the infrastructure domain where protected attributes may be spatial rather than individual and where the definition of equitable outcomes is contested. The governance challenge is to define acceptable fairness criteria through inclusive stakeholder engagement and to embed fairness auditing into the lifecycle of reliability analytics, a practice that has not yet been institutionalized in civil infrastructure asset management.

The sustainability of fault detection infrastructure itself merits critical examination. Pervasive sensing requires the production and deployment of millions of sensor nodes, each containing rare earth elements, batteries, and electronic components whose manufacturing and disposal carry substantial environmental footprints. Edge gateways and cloud data centers consume significant electricity, and continuous wireless data transmission depletes batteries that must be periodically replaced, particularly in remote infrastructure assets such as pipeline monitoring stations in deserts or offshore platforms. Life cycle assessment methodologies [21] can be applied to quantify the environmental burden of the data infrastructure relative to the reliability benefits gained, such as the reduction in catastrophic failures that release pollutants or consume large quantities of energy for emergency repairs. However, such holistic assessments are rarely performed before large-scale sensor rollouts, leading to a situation where the environmental costs of digitalization are externalized. A sustainable approach to data-enabled reliability requires designing sensor systems for low power, energy harvesting, recyclability, and modular repairability, as well as policies that internalize the life-cycle environmental costs of the digital infrastructure layer.

Adversarial robustness introduces another dimension of sustainability and trust. As fault detection models become integrated into safety-critical decision pipelines, they become attractive targets for adversaries seeking to disrupt infrastructure services through data poisoning during training or spoofing attacks during inference [20]. Injection of carefully crafted sensor perturbations can, for example, cause a wind turbine gearbox fault detector to suppress alerts while damage accumulates, leading to costly unscheduled outages or safety incidents. Defenses such as adversarial training, input sanitization, and sensor redundancy add complexity and must be weighed against the benefits. The long-term sustainability of trust in

automated infrastructure monitoring hinges on the ability to maintain robustness in the face of an evolving threat landscape, which necessitates continuous investment in cybersecurity research, vulnerability disclosure programs, and international coordination on critical infrastructure protection norms.

6. Policy Implications and Future Research Trajectories

The policy landscape for data-enabled fault detection and reliability assessment is currently fragmented, with different regulatory regimes governing the physical asset, the communication infrastructure, and the algorithmic decision-making processes. In the United States, for instance, the electric grid is subject to mandatory reliability standards enforced by the North American Electric Reliability Corporation, while the algorithmic condition monitoring tools used by utilities are largely unregulated as long as they do not replace mandatory protection functions. This regulatory asymmetry creates a gap in which infrastructure operators may deploy machine learning-based predictive maintenance with inadequate oversight, while regulatory bodies lack the in-house technical expertise to audit model risk. A coherent policy framework must establish performance-based criteria for data-enabled fault detection systems that are proportional to the safety or societal consequences of misdiagnosis. Such criteria should require operators to demonstrate model validation on representative operational datasets, ongoing drift monitoring, and traceability of decisions that affect asset replacement or safety-critical control actions.

Certification of learning-enabled components in infrastructure invites comparison with the aviation industry's development assurance processes for complex electronic hardware and software. However, the continuous learning paradigm, where models are retrained periodically on newly accumulated data, challenges the traditional certification model that assumes a frozen design baseline. Policy innovations are needed to define "safe evolutions" of data-driven models, perhaps through performance envelopes within which model updates can be deployed without full recertification, combined with automated regression testing suites that capture domain knowledge about unacceptable failure modes. Public-private partnerships could sponsor shared testing facilities and reference datasets that allow independent benchmarking of fault detection algorithms across vendors, analogous to the role of the National Institute of Standards and Technology in facial recognition evaluation, but tailored to infrastructure failure modes.

A second policy frontier concerns data sharing across infrastructure sectors and organizational boundaries. Fault signatures learned from one water utility's pump fleet could accelerate diagnostics for another utility, but proprietary concerns and liability fears inhibit data pooling. Federated learning offers a technical path toward collaborative model training without centralizing raw data, yet governance structures for federated model ownership, contribution accounting, and benefit-sharing are nascent. Policymakers could incentivize data collaboratives through safe harbor provisions that limit liability for shared data used for safety improvement, while mandating strong data anonymization and cybersecurity protections. Interoperability mandates for data formats and application programming interfaces, modeled on open banking regulations, could reduce switching costs and prevent vendor lock-in, thereby fostering a competitive ecosystem of analytics providers and accelerating innovation cycles in condition monitoring technologies.

Future research must bridge the silos between reliability engineering, machine learning operations, human factors, and public policy. Longitudinal field studies that track the performance of data-enabled fault detection systems over multiple maintenance cycles are

urgently needed to empirically ground claims of reliability improvement and to quantify the hidden costs of model maintenance. Developing causal inference methods that can distinguish correlation-induced spurious fault indicators from true degradation causation will be essential to avoid unnecessary interventions that waste resources and undermine trust. Research on explainable artificial intelligence for temporal sensor data should move beyond saliency maps to produce explanations that align with the mental models and decision workflows of maintenance engineers and reliability analysts. Finally, the sustainability of the digital infrastructure itself must become a core research theme, requiring cross-disciplinary studies that combine life cycle assessment, energy-aware computing, and infrastructure policy to chart pathways toward circular economy models for sensor networks. The integration of these research threads will determine whether data-enabled fault detection evolves into a robust, equitable, and enduring pillar of modern infrastructure stewardship.

7. Conclusion

Data-enabled fault detection and reliability assessment represent a transformative capability for modern engineering infrastructure, yet their full potential remains contingent on addressing deep structural challenges that span architecture, governance, human factors, policy, and environmental sustainability. The preceding analysis has argued that the reliability of a data-driven monitoring system cannot be evaluated in isolation from the sociotechnical infrastructure that sustains it, including data pipelines, model lifecycle management, institutional accountability structures, and regulatory frameworks. Central architectural trade-offs between edge and cloud processing, between physical model fidelity and data-driven scalability, and between model accuracy and explainability must be navigated with an awareness of their long-term implications for maintainability, auditability, and resilience to distributional shift. Governance gaps in data ownership, model certification, and fairness persist and demand new institutional arrangements that involve utilities, technology vendors, regulators, and the communities served by infrastructure. Sustainability concerns add urgency to the design of resource-efficient sensing and computing substrates and to the incorporation of life-cycle thinking into digitalization strategies. Policy and research communities must collaborate to develop performance-based regulatory frameworks, shared validation infrastructure, and incentives for data sharing that accelerate learning while protecting public safety and equity. A system-level perspective, attentive to both technical and social dimensions, is essential to ensure that the data-enabled transformation of infrastructure reliability delivers on its promise without creating new fragilities or deepening existing inequities. The path forward requires interdisciplinary scholarship that integrates reliability science, machine learning, human-centered design, and policy analysis into a coherent discipline of data-enabled infrastructure stewardship.

References

1. Zio, E. (2016). Challenges in the vulnerability and risk analysis of critical infrastructures. *Reliability Engineering & System Safety*, 152, 137–150. <https://doi.org/10.1016/j.ress.2016.02.009>
2. Jardine, A. K. S., Lin, D., & Banjevic, D. (2006). A review on machinery diagnostics and prognostics implementing condition-based maintenance. *Mechanical Systems and Signal Processing*, 20(7), 1483–1510. <https://doi.org/10.1016/j.ymssp.2005.09.012>
3. Rausand, M., & Høyland, A. (2004). *System reliability theory: Models, statistical methods, and applications* (2nd ed.). John Wiley & Sons.

4. Tao, F., Zhang, H., Liu, A., & Nee, A. Y. C. (2019). Digital twin in industry: State-of-the-art. *IEEE Transactions on Industrial Informatics*, 15(4), 2405–2415.
<https://doi.org/10.1109/TII.2018.2873186>
5. 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>
6. Montgomery, D. C. (2007). *Introduction to statistical quality control* (5th ed.). John Wiley & Sons.
7. Isermann, R. (2005). Model-based fault-detection and diagnosis – status and applications. *Annual Reviews in Control*, 29(1), 71–85.
<https://doi.org/10.1016/j.arcontrol.2004.12.002>
8. Lei, Y., Yang, B., Jiang, X., Jia, F., Li, N., & Nandi, A. K. (2020). Applications of machine learning to machine fault diagnosis: A review and roadmap. *Mechanical Systems and Signal Processing*, 138, 106587. <https://doi.org/10.1016/j.ymssp.2019.106587>
9. Chandola, V., Banerjee, A., & Kumar, V. (2009). Anomaly detection: A survey. *ACM Computing Surveys*, 41(3), 1–58. <https://doi.org/10.1145/1541880.1541882>
10. Venkatasubramanian, V., Rengaswamy, R., Yin, K., & Kavuri, S. N. (2003). A review of process fault detection and diagnosis: Part I: Quantitative model-based methods. *Computers & Chemical Engineering*, 27(3), 293–311. [https://doi.org/10.1016/S0098-1354\(02\)00160-6](https://doi.org/10.1016/S0098-1354(02)00160-6)
11. Stamatelatos, M., Dezfuli, H., Apostolakis, G., Everline, C., Guarro, S., Mathias, D., Mosleh, A., & Youngblood, R. (2011). *Probabilistic risk assessment procedures guide for NASA managers and practitioners* (NASA/SP-2011-3421). NASA.
12. Weber, P., Medina-Oliva, G., Simon, C., & Iung, B. (2012). Overview on Bayesian networks applications for dependability, risk analysis and maintenance areas. *Engineering Applications of Artificial Intelligence*, 25(4), 671–682.
<https://doi.org/10.1016/j.engappai.2010.06.002>
13. Labeau, P. E., Smidts, C., & Swaminathan, S. (2000). Dynamic reliability: Towards an integrated platform for probabilistic risk assessment. *Reliability Engineering & System Safety*, 68(3), 219–254. [https://doi.org/10.1016/S0951-8320\(00\)00004-9](https://doi.org/10.1016/S0951-8320(00)00004-9)
14. Francis, R., & Bekera, B. (2014). A metric and frameworks for resilience analysis of engineered and infrastructure systems. *Reliability Engineering & System Safety*, 121, 90–103. <https://doi.org/10.1016/j.ress.2013.07.004>
15. Shi, W., Cao, J., Zhang, Q., Li, Y., & Xu, L. (2016). Edge computing: Vision and challenges. *IEEE Internet of Things Journal*, 3(5), 637–646.
<https://doi.org/10.1109/JIOT.2016.2579198>
16. Khatri, V., & Brown, C. V. (2010). Designing data governance. *Communications of the ACM*, 53(1), 148–152. <https://doi.org/10.1145/1629175.1629210>
17. ISO 13374-1:2003. *Condition monitoring and diagnostics of machines — Data processing, communication and presentation — Part 1: General guidelines*. International Organization for Standardization.

18. De Bruijn, H., & Herder, P. M. (2009). System and actor perspectives on sociotechnical systems. *IEEE Transactions on Systems, Man, and Cybernetics - Part A: Systems and Humans*, 39(5), 981–992. <https://doi.org/10.1109/TSMCA.2009.2025452>
19. 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), 1–35. <https://doi.org/10.1145/3457607>
20. Goodfellow, I. J., Shlens, J., & Szegedy, C. (2015). Explaining and harnessing adversarial examples. In *International Conference on Learning Representations (ICLR)*. arXiv:1412.6572.
21. Hauschild, M. Z., Rosenbaum, R. K., & Olsen, S. I. (Eds.). (2018). *Life cycle assessment: Theory and practice*. Springer. <https://doi.org/10.1007/978-3-319-56475-3>