Operational failure or error?
Enforcing effective accident investigation

Avi Harel

Ergolight

October 2025

Abstract

The article compares two strategies for accident investigation: traditional Root Cause Analysis (RCA) and Multi-Layer Cause Analysis (MLCA) commonly employed in industries of high-risk operation. RCA tends to produce blame-oriented outcomes, while MLCA facilitates systemic learning. The comparison is in terms of learning from accidents and opportunities for improving. Blaming human operators often masks latent design flaws, organizational pressures, and insufficient safety culture. Drawing on case studies from aviation, nuclear energy, healthcare, and industrial automation, the article illustrates how superficial safety culture leads organizations to focus on human error rather than systemic weaknesses. In the context of the Fifth Industrial Revolution (5IR), where cyber-physical systems demand integrated resilience, MLCA offers a more sustainable framework for preventing repeated failures and promoting organizational antifragility. The article discusses the effect of safety culture on the effectiveness of RCA and proposes that in the 5IR the MLCA should dominate the RCA strategy.

Keywords

Operational failure, human error, safety culture, Root Cause Analysis (RCA), Multi-Layer Cause Analysis (MLCA), Fifth Industrial Revolution (5IR), operational sustainability, antifragility

Introduction

Operational failures are exceptional events that produce unintended, costly, and often hazardous outcomes in complex socio-technical systems. These failures typically emerge from the interaction of human, technical, and organizational components. Yet, in the aftermath of such events, industry investigations and media narratives frequently attribute the cause to “operator error.” This default explanation simplifies complex causal chains and satisfies the organizational and public desire for a clear, actionable answer. However, such attributions often obscure systemic deficiencies in design, management, or safety culture.

Traditionally, people regard the operators as responsible for preventing failure, and when they fail, industry investigations and media narratives frequently attribute the cause to 'operator error.' Donald Norman (2016) famously critiques the reflexive “human error” attribution:

> “Over 90% of industrial accidents are blamed on human error. You know, if it was 5%, we might believe it. But when it is virtually always, shouldn't we realize that it is something else? “
>
> “When there are serious accidents, the first reaction is often to claim ‘human error.’ … We won’t solve these problems until we recognize that bad design of equipment and procedures is most often the culprit.”

This default explanation simplifies complex causal chains and satisfies the organizational and public desire for a clear, actionable answer. Human error is often misattributed and that system design is usually the true source of failure.

> “This is a common result when design is done by engineers who assume logical functioning and have not themselves studied the complexities of operational behavior, with multiple people, continual interruptions, and unexpected situations.”

Norman’s critique reframes operator error as a symptom of poor system design rather than a cause of accidents.

Much of what is called “human error” is better understood as a breakdown of design, context, or system constraints. In effect, blaming the operator is a default, low-effort response that avoids questioning deeper structural causes.

This aligns with a broader systems thinking perspective that views accidents as the result of multiple interacting failures across technical, procedural, and organizational layers. Despite this, the persistence of “human error” narratives suggests that organizational and psychological factors, particularly those linked to safety culture, shape how failures are interpreted and reported.

The industry-wide tendency to attribute failures to human error carries significant costs. When organizations focus on individual blame, they forfeit opportunities for organizational learning, resilience building, and antifragility; the capacity to improve through stress and failure (Taleb, 2012). Instead of analyzing the latent conditions that allowed the failure to occur, such organizations treat the incident as an anomaly caused by a “bad apple” (Dekker, 2011). The result is a cycle of recurring vulnerabilities, masked by the illusion of corrective action.

This paper investigates the question: Why are operational failures often attributed to operator errors? It proposes that this phenomenon arises not from empirical evidence of

individual fault but from organizational dynamics, specifically, the tension between short-term and long-term goals, the depth of safety culture implementation, and the strategic framing of investigations.

To examine this question, the article contrasts two dominant investigative strategies, namely, Root Cause Analysis (RCA) and Multi-Layer Cause Analysis (MLCA) and explores how their use reflects the maturity of an organization's safety culture. It argues that the dominance of RCA in many industrial domains is symptomatic of shallow safety culture and that, in the context of the Fifth Industrial Revolution (5IR), MLCA should replace RCA as the preferred approach.

Operational Failures as Operator Error

Why are operational failures attributed to operator error?

Attributing operational failures to operator errors is rarely a neutral or purely technical act. Rather, it reflects how organizations interpret events in light of their goals, culture, and stakeholder pressures. This section examines the organizational dynamics that shape such attribution and explains why the tendency to blame operators is most pronounced in organizations with superficial safety cultures.

Organizational Goals and Investigation Target

Organizations operate under competing imperatives. On the one hand, they pursue **long-term goals**: preventing recurrence, building resilience, and sustaining safety. On the other hand, they must satisfy short-term goals: reducing operational costs, maintaining productivity, and demonstrating accountability to external stakeholders.

When failures occur, the tension between these two goal sets becomes acute. In a mature safety culture, management prioritizes long-term learning: understanding latent conditions, redesigning flawed processes, and strengthening organizational defenses (Reason, 1997). However, when safety culture is superficial, management tends to prioritize short-term optics. The investigation thus becomes a public performance of control rather than an authentic search for causes (Hopkins, 2019).

In such cases, blaming the operator is expedient. The operator is visible, substitutable, and symbolically culpable - qualities that make human error the "perfect scapegoat." This attribution protects organizational reputation, reassures stakeholders, and minimizes the costs of systemic reform.

The Role of Safety Culture

Safety culture refers to shared values, beliefs, and behaviors that determine how an organization prioritizes safety relative to productivity or profitability (Reason, 1997; Guldenmund, 2010). When implemented properly, safety culture promotes transparency, learning, and accountability. It encourages reporting of near misses and systemic vulnerabilities.

Conversely, a superficial safety culture uses the language of safety without embedding it into daily decision-making. Indicators of superficiality include:

- Low tolerance for uncertainty or failure
- Reliance on punishment rather than learning
- Minimal cross-functional learning from incidents, and
- Token compliance with safety audits.

Organizations with such cultures are prone to cognitive simplification, reducing complex systemic issues to singular, human-centered causes.

Management–Investigator Alignment

The alignment between management goals and investigative scope strongly influences the outcome of accident investigations. Management typically appoints investigators who understand the operational context and who are trusted to protect institutional interests. When long-term safety culture is robust, this alignment yields transparent inquiry and system-wide learning.

However, in organizations with weak safety culture, investigators operate under implicit pressure to produce "acceptable" conclusions, findings that demonstrate managerial control without exposing systemic design flaws. This often results in confirmation bias: investigators unconsciously seek evidence supporting the pre-existing narrative of human error.

While blaming an operator may satisfy immediate reputational needs, it imposes hidden long-term costs. It discourages reporting, suppresses learning, and leaves latent system vulnerabilities unaddressed. As Reason (2008) observed, "blame is emotionally satisfying but operationally futile." Organizations that substitute accountability for understanding inadvertently create the conditions for recurrence.

When safety culture is **deeply implemented**, organizations align their long-term and short-term goals, leading to MLCA-style investigations and sustained improvement. Conversely, **superficial implementation** of safety culture leads to RCA-focused investigations that privilege short-term optics and human error attribution.

Investigation Strategies: RCA vs. MLCA

Accident investigation is both a technical and cultural process. The selected investigative method reflects not only the complexity of the system but also the organization's values, priorities, and tolerance for uncertainty. Two dominant strategies are used in industry today: Root Cause Analysis (RCA) and Multi-Layer Cause Analysis (MLCA).

RCA seeks a single root cause, whereas MLCA examines multiple layers of failure. The choice of strategy reflects both technical and cultural perspectives. While RCA seeks a single underlying "root" cause that, once corrected, prevents recurrence, MLCA explores how multiple protective layers failed simultaneously. The distinction between these methods illustrates the gap between simplified blame attribution and systemic learning.

Root Cause Analysis (RCA)

Root Cause Analysis emerged from engineering and quality management traditions (e.g., Deming, 1986; Imai, 1986). It aims to identify the most fundamental cause whose elimination prevents recurrence. RCA tools include the Five Whys, fishbone (Ishikawa) diagrams, fault tree analysis, and event and causal factor charts.

Although widely used in healthcare, manufacturing, and industrial safety, RCA has been criticized for its linearity and its implicit assumption that accidents have a single identifiable cause (Peerally et al., 2017). Such simplicity appeals to management because it produces closure and enables direct accountability. However, it fails to account for the interactive complexity and systemic coupling that characterize modern socio-technical systems (Perrow, 1999).

Peerally et al. (2017) identified eight recurring challenges with RCA in healthcare: oversimplification, hindsight bias, confirmation bias, limited follow-up, poor system representation, narrow scope, lack of interdisciplinary integration, and failure to generate learning. These weaknesses mirror issues seen in other domains, from industrial process safety to aviation.

Multi-Layer Cause Analysis (MLCA)

In contrast, Multi-Layer Cause Analysis extends James Reason's Swiss cheese model (1997), which conceptualizes systems as comprising multiple layers of defense—technological, procedural, and organizational. Each layer contains "holes" (latent weaknesses) that can align to permit failure.

MLCA does not seek a single root cause but analyzes how defenses eroded, how barriers interacted, and how human, organizational, and technical elements co-produced the event. This method views accidents as systemic breakdowns rather than discrete errors. It encourages resilience engineering and organizational learning (Hollnagel et al., 2006).

In MLCA, each layer is examined to determine:

- What defenses failed and why
- How the system responded under stress, and
- What modifications could strengthen cross-layer resilience.

This approach is now standard in aviation, nuclear, and aerospace safety investigations (Leveson, 2011; Dekker, 2011). These sectors recognize that eliminating "human error" is neither possible nor desirable; instead, the goal is to understand system conditions that enable or constrain human performance.

Comparison

The choice between RCA and MLCA is not purely methodological; it signals an organization's epistemological stance toward failure. RCA aligns with a control-oriented worldview: failures are anomalies caused by deviations from rules. MLCA aligns with a complexity-oriented worldview: failures are emergent, requiring systemic understanding. The following table compares of RCA with MLCA, based on data from Reason (1997), Peerally et al. (2017), Hollnagel et al. (2006), and Leveson (2011).

Table 1. Comparison of Root Cause Analysis (RCA) and Multi-Layer Cause Analysis (MLCA)

Dimension	Root Cause Analysis (RCA)	Multi-Layer Cause Analysis (MLCA)
Philosophical basis	Linear causality; seeks a single root cause	Systemic causality; examines multiple interacting causes
Primary question	"Who or what failed?"	"How did multiple barriers fail together?"
Causal model	Sequential event chain	Layered defense and interaction model
Focus	Error identification and correction	System resilience and learning
Typical tools	Five Whys, Ishikawa diagram, fault tree	Swiss cheese model, barrier analysis, STAMP
Treatment of human error	Viewed as primary cause	Viewed as system symptom or adaptation
Organizational implication	Assigns blame or accountability	Encourages shared responsibility and learning
Outcome orientation	Quick corrective actions	Long-term system redesign and safety culture improvement
Common application domains	Manufacturing, healthcare, maintenance	Aviation, nuclear, aerospace, high-risk industries
Alignment with safety culture maturity	Low to moderate	High (learning-oriented, non-punitive)

Methodological and Cultural Implications

When safety culture is mature, MLCA dominates because the organization recognizes that complexity demands multi-layer learning. When safety culture is superficial, RCA prevails because it offers a quick, politically safe explanation, usually pointing to operator error.

Figure 1. Organizational Attribution Model for Operational Failures

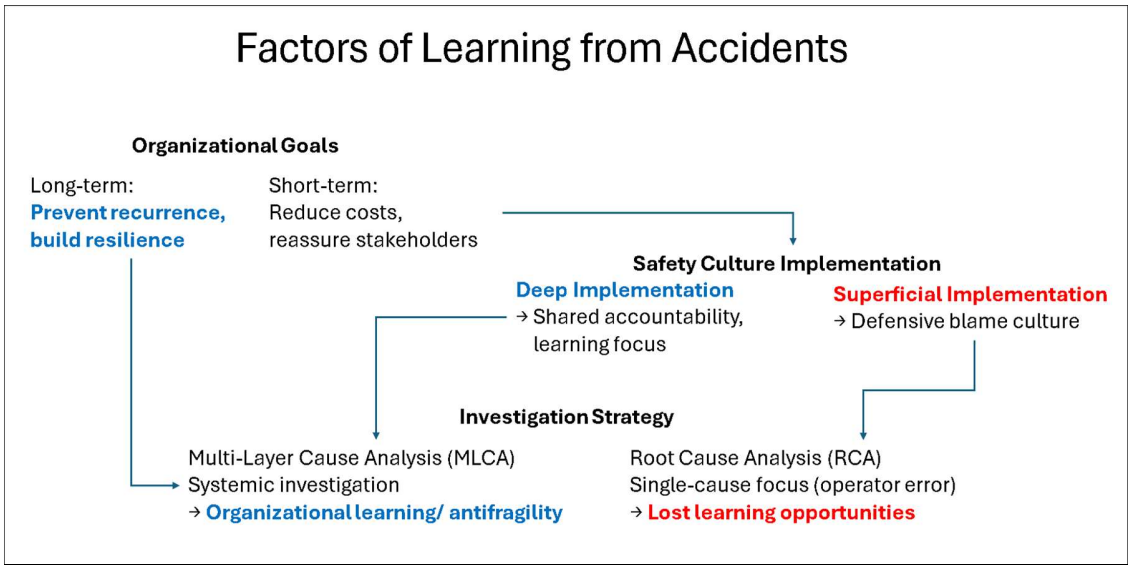

In practice, organizations often combine both methods: RCA provides an initial causal pathway, while MLCA expands it into a systemic map of vulnerabilities. The challenge lies in ensuring that the deeper layers of analysis are not abandoned once a convenient human error explanation is found.

Case Studies of Operational Failure Attribution

Real-world accidents reveal how the attribution of operational failures depends on the maturity of safety culture and the chosen investigation strategy. This section examines representative cases from several industrial domains: aviation, energy, healthcare, and manufacturing, to illustrate how the attribution of operational failures to operator error reflects underlying differences in safety culture maturity.

Aviation: The Tenerife Airport Disaster

Year: 1977
Initial Attribution: Pilot error
Investigation Type: Transition to MLCA (CRM)

Safety Culture Maturity: Emerging
Key Lesson: Communication and teamwork failures reflect system design, not individuals.

Aviation: Air France Flight 447

Year: 2009
Initial Attribution: Pilot confusion
Investigation Type: MLCA (BEA)
Safety Culture Maturity: Mature
Key Lesson: Systemic view improves training, design, and resilience.

Nuclear Energy: Three Mile Island

Year: 1979
Initial Attribution: Operator misjudgment
Investigation Type: MLCA (Defense-in-Depth)
Safety Culture Maturity: Developing
Key Lesson: Layered safety barriers essential; design must anticipate error.

Nuclear Energy: Chernobyl

Year: 1986
Initial Attribution: Operator negligence
Investigation Type: RCA (political bias)
Safety Culture Maturity: Weak
Key Lesson: Poor safety culture promotes scapegoating and repeated failure.

Healthcare: Surgical and Diagnostic Errors

Year: 2000s
Initial Attribution: Clinician error
Investigation Type: MLCA (Systemic Safety)
Safety Culture Maturity: Developing
Key Lesson: System redesign prevents recurring medical errors.

Industrial: Texas City & Deepwater Horizon

Year: 2005–2010
Initial Attribution: Operator/contractor error
Investigation Type: MLCA (CSB Analysis)
Safety Culture Maturity: Weak to moderate
Key Lesson: Cost-driven RCA conceals latent conditions; MLCA promotes resilience.

Summary of case studies

Table 2. Summary of Case Studies: Attribution Type, Investigation Strategy, and Lessons Learned

Case / Sector	Year	Initial Attribution	Investigation Type	Safety Culture Maturity	Key Lesson
Aviation: The Tenerife Airport Disaster	1977	Pilot error	Transition to MLCA (CRM)	Emerging	Communication and teamwork failures reflect system design, not individuals.
Aviation: Air France Flight 447	2009	Pilot confusion	MLCA (BEA)	Mature	Systemic view improves training, design, and resilience.
Nuclear Energy: Three Mile Island	1979	Operator misjudgment	MLCA (Defense-in-Depth)	Developing	Layered safety barriers essential; design must anticipate error.
Nuclear Energy: Chernobyl	1986	Operator negligence	RCA (political bias)	Weak	Poor safety culture promotes scapegoating and repeated failure.
Healthcare: Surgical and Diagnostic Errors	2000s	Clinician error	MLCA (Systemic Safety)	Developing	System redesign prevents recurring medical errors.
Industrial: Texas City & Deepwater Horizon	2005–2010	Operator/contractor error	MLCA (CSB Analysis)	Weak to moderate	Cost-driven RCA conceals latent conditions; MLCA promotes resilience.

Discussion of Patterns

Across sectors, a consistent pattern emerges:

- Superficial safety cultures correlate with RCA and operator blame.
- Mature safety cultures employ MLCA, focusing on multi-layer resilience.
- Transitioning to MLCA yields antifragile learning (Taleb, 2012); organizations not only recover from failures but adapt and strengthen.

These findings support the argument that operational failures are often attributed to operator error not because operators are inherently at fault, but because organizations lacking deep safety culture seek short-term legitimacy through blame rather than long-term improvement.

Implications for the Fifth Industrial Revolution (5IR)

Operational failures are often attributed to operator error not because human operators are inherently unreliable, but because organizations prioritize short-term optics over long-term learning. In the context of the Fifth Industrial Revolution, this paradigm is obsolete.

From Blame to Learning: The 5IR Shift

The Fifth Industrial Revolution (5IR) marks a paradigm shift in how humans and intelligent systems collaborate. Unlike the Fourth Industrial Revolution (4IR), which emphasized automation, cyber-physical systems, and efficiency the 5IR emphasizes collaboration, ethics, and sustainability (Xu et al., 2021). Within this context, human–system integration (HSI) and safety culture converge toward a single goal: sustainable operational learning.

Blaming operators for failures is incompatible with the principles of 5IR. Such blame not only discourages learning but also contradicts the antifragility concept proposed by Taleb (2012), the idea that systems should gain from disorder rather than merely resist it. Instead of masking deficiencies through scapegoating, 5IR-oriented organizations leverage operational failures as feedback mechanisms to strengthen both human and artificial agents within socio-technical systems.

MLCA as a Model for 5IR Safety Culture

The Multi-Layer Cause Analysis (MLCA) approach aligns closely with 5IR values because it emphasizes:

- Systemic understanding: Viewing accidents as emergent from complex interactions rather than isolated human mistakes.
- Collaborative intelligence: Integrating human insight and AI-driven diagnostics to uncover deep causal networks.
- Sustainable improvement: Translating failures into design feedback loops that enhance resilience over time.

In a 5IR environment, safety analysis becomes a cyber-human partnership. AI systems can assist by identifying patterns across thousands of incident reports, while human investigators contextualize these insights within operational realities. Together, they form an antifragile learning ecosystem, capable of adapting dynamically to emerging risks.

Antifragility and Organizational Learning

Taleb (2012) argued that traditional risk management aims to minimize volatility, whereas antifragile systems benefit from it. This distinction applies directly to operational safety.

- Fragile systems—those driven by short-term goals and RCA mindsets—crack under stress because they conceal latent vulnerabilities.
- Resilient systems absorb shocks but return merely to their previous state.
- Antifragile systems, however, evolve by learning from every disruption.

MLCA embodies antifragility: each identified layer of defense (technical, human, organizational) becomes a node of potential improvement. Instead of seeking the one "root cause," antifragile organizations view incidents as evolutionary catalysts for innovation and adaptation.

Figure 2. The Antifragile Learning Loop in 5IR Operations

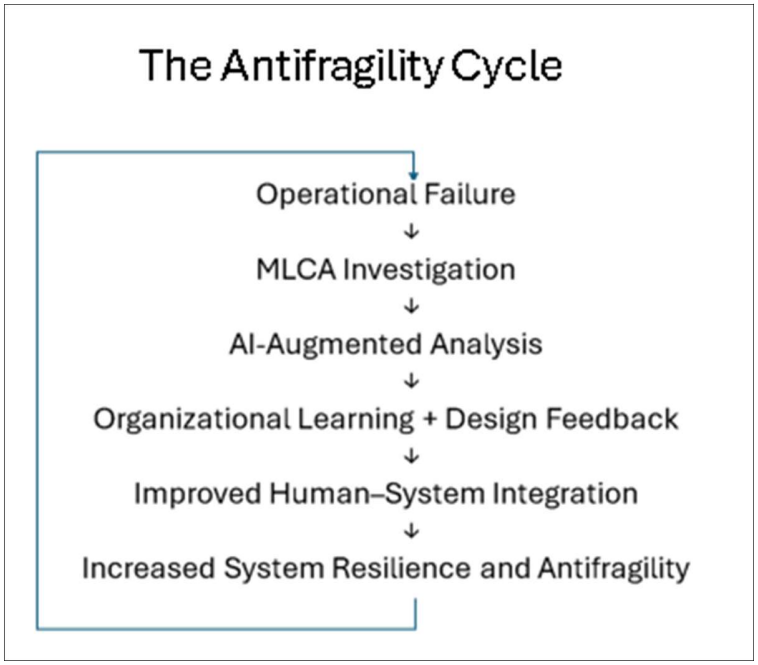

Integrating Generative AI into MLCA

Generative AI (e.g., large language models) enhances MLCA through four mechanisms:

- Pattern detection: Identifying recurrent causal combinations across large incident datasets.
- Scenario modeling: Simulating how latent conditions could interact under varying circumstances.
- Knowledge synthesis: Linking technical documentation, human reports, and historical cases into unified causal maps.
- Real-time decision support: Assisting operators during abnormal situations by suggesting context-aware responses.

These capabilities transform the investigation process from a retrospective blame exercise into a prospective learning mechanism—a cornerstone of 5IR safety philosophy (Ramos et al., 2023).

Implications for Policy and Practice

By adopting Multi-Layer Cause Analysis, organizations can shift from a culture of blame to one of antifragile learning, integrating generative AI and human expertise to create systems that thrive under uncertainty. This represents not only a technical evolution but a moral and organizational transformation—placing sustainability, responsibility, and learning at the heart of modern operations.

To operationalize the 5IR vision, organizations may:

- Institutionalize MLCA as the standard investigative framework, replacing RCA in all but the simplest cases.
- Reframe safety culture as a system-wide competence integrating human judgment, AI analytics, and ethical governance.
- Develop antifragile metrics: measuring improvement from incidents, not just avoidance of harm.
- Foster transparency and trust between human and artificial agents, ensuring accountability without blame.
- Invest in learning infrastructure that captures, analyzes, and disseminates lessons from operational experiences.

These measures align industrial practices with both **humanistic ethics** and **technological evolution**, ensuring that operational systems remain not only safe but also continuously improving.

Conclusions

Operational failures are often attributed to operator errors because organizations prioritize short-term optics. Operational failures are often attributed to operator error not because human operators are inherently unreliable, but because organizations prioritize short-term optics over long-term learning. In the context of the Fifth Industrial Revolution, this paradigm is obsolete.

The Fifth Industrial Revolution requires organizations to adopt MLCA, deep safety culture, and antifragile learning loops. By adopting Multi-Layer Cause Analysis, organizations can shift from a culture of blame to one of antifragile learning, integrating generative AI and human expertise to create systems that thrive under uncertainty.

This represents not only a technical evolution but a moral and organizational transformation, placing sustainability, responsibility, and learning at the heart of modern operations. Practical steps include embedding MLCA, leveraging AI, and aligning management goals with long-term learning and resilience.

Key Findings

This study has examined why operational failures are often attributed to operator error, drawing on literature, investigation strategies, and real-world case studies across aviation, nuclear energy, healthcare, and industrial domains. The key findings are:

- Operator blame is frequently a symptom, not the cause: Investigations that focus on human error often overlook latent organizational, technical, and procedural deficiencies (Norman, 2016; Dekker, 2011).
- Safety culture drives investigation strategy: Superficial safety cultures favor Root Cause Analysis (RCA) and short-term optics, producing human-centric attributions. Mature safety cultures favor Multi-Layer Cause Analysis (MLCA), encouraging systemic learning and resilience (Reason, 1997; Hollnagel et al., 2006).
- Multi-layer analysis prevents recurrence: Across sectors, MLCA identifies interacting failures in technical, human, and organizational layers. This approach supports long-term organizational learning and aligns with antifragile principles (Taleb, 2012).
- 5IR requires a paradigm shift: In the Fifth Industrial Revolution, human–system integration (HSI) and AI-assisted analysis make operator-blame strategies obsolete. Failures should instead feed into adaptive learning loops, improving system performance over time.

Practical Recommendations for Organizations (Narrative Version)

To effectively address operational failures in the context of modern socio-technical systems, organizations must integrate safety culture, investigative rigor, and adaptive learning into their operational framework. First, Multi-Layer Cause Analysis (MLCA) should be institutionalized as the standard approach for investigating incidents. Unlike traditional Root Cause Analysis (RCA), which tends to isolate a single cause, MLCA examines interactions across multiple technical, human, and organizational layers, enabling a more comprehensive understanding of failure dynamics.

Embedding a deep safety culture is equally essential. Organizations must promote transparency, open reporting, and collective learning at all levels. This involves cultivating an environment where employees and managers alike feel empowered to identify latent risks without fear of punitive consequences. Such a culture encourages collaboration across technical, managerial, and operational units, reinforcing the systemic view of safety.

In parallel, leveraging advanced AI and data analytics can enhance organizational learning. Generative AI models and predictive analytics tools allow organizations to detect recurring patterns, anticipate latent risks, and support decision-making in real time. By integrating AI with human expertise, organizations create a synergistic environment where insights from operational data translate directly into actionable improvements.

Antifragile feedback loops should be implemented to transform incidents into opportunities for growth. Rather than simply returning systems to their previous state, these loops use failures as a source of learning that strengthens both human operators and technical systems. This approach ensures that each incident contributes to increasing system resilience over time.

Finally, management goals must align with long-term learning objectives rather than short-term optics. Organizations should prioritize sustainable improvements and system-wide resilience over immediate cost reduction or public relations considerations. Human operators should be treated as co-designers of the system, with training and decision-making authority that reflects their integral role in maintaining safety. By embracing these measures, organizations can shift from reactive, blame-oriented approaches to proactive, learning-centered strategies that fully realize the principles of the Fifth Industrial Revolution.

Closing Remarks (Narrative Version)

Operational failures reveal the true maturity of an organization's safety culture. When management chooses to attribute breakdowns to individual operator errors, it exposes a defensive mindset that values short-term reassurance over genuine system learning. Conversely, when an organization embraces Multi-Layer Cause Analysis (MLCA) and antifragile thinking, each failure becomes a catalyst for progress rather than a signal of blame.

In the emerging landscape of the Fifth Industrial Revolution (5IR), the distinction between human and artificial agents is increasingly blurred. Operators, engineers, and intelligent systems must collaborate within adaptive frameworks that learn continuously from experience. The organizations that thrive will be those that recognize human fallibility not as a weakness to conceal, but as a critical source of information about how complex systems behave in practice.

True resilience arises when failure is treated as feedback. By integrating antifragile principles, fostering psychological safety, and aligning managerial intent with long-term organizational learning, institutions can evolve beyond traditional safety paradigms. In this model, human error is reframed as design information; investigation becomes a learning process; and safety culture transforms from a compliance requirement into a strategic advantage.

As industries enter the 5IR, the capacity to learn adaptively—through MLCA, AI-supported analysis, and open reporting—will define operational excellence. The challenge is not to eliminate error but to learn faster and more effectively from it. When organizations achieve that, operational failures cease to be endpoints and become essential drivers of sustainable innovation.

References

Bureau d'Enquêtes et d'Analyses (BEA). (2012). Final report on the accident on 1st June 2009 to the Airbus A330-203 registered F-GZCP operated by Air France flight AF447. Paris: BEA.

Chemical Safety Board (CSB). (2016). BP Texas City refinery explosion investigation report. U.S. Chemical Safety Board.

Dekker, S. (2011). Drift into failure: From hunting broken components to understanding complex systems. Ashgate.

Deming, W. E. (1986). Out of the crisis. MIT Press.

Guldenmund, F. W. (2010). (Mis)understanding safety culture and its relationship to safety management. Safety Science, 48(8), 1087–1092. https://doi.org/10.1016/j.ssci.2009.10.010

Hopkins, A. (2019). Learning from disasters: A managerial approach. CCH Australia.

Imai, M. (1986). Kaizen: The key to Japan's competitive success. McGraw-Hill.

Leveson, N. (2011). Engineering a safer world: Systems thinking applied to safety. MIT Press.

Norman, D. A. (2016). Human error? No, bad design. JND.org. https://jnd.org/human-error-no-bad-design/

Peerally, M. F., Carr, S., Waring, J., & Dixon-Woods, M. (2017). The problem with root cause analysis. BMJ Quality & Safety, 26(5), 417–422. https://doi.org/10.1136/bmjqs-2016-005511

Perrow, C. (1999). *Normal Accidents: Living with High-Risk Technologies* (expanded ed.). Princeton University Press.

Reason, J. (1997). *Managing the risks of organizational accidents.* Aldershot, UK: Ashgate.

Reason, J. (2008). The human contribution: Unsafe acts, accidents and heroic recoveries. CRC Press.

Ramos, C., Cordeiro, J., & Pereira, A. (2023). AI-supported safety management in socio-technical systems: Toward adaptive learning frameworks. Safety Science, 167, 106187. https://doi.org/10.1016/j.ssci.2023.106187

Taleb, N. N. (2012). Antifragile: Things that gain from disorder. Random House.

Xu, M., David, J. M., & Kim, S. H. (2021). The Fourth Industrial Revolution: Opportunities and challenges. International Journal of Financial Research, 12(3), 90–104. https://doi.org/10.5430/ijfr.v12n3p90