Proceedings of the 35th European Safety and Reliability & the 33rd Society for Risk Analysis Europe Conference
Edited by Eirik Bjorheim Abrahamsen, Terje Aven, Frederic Bouder, Roger Flage, Marja Ylönen

doi: 10.3850/978-981-94-3281-3_ESREL-SRA-E2025-P0487-cd

Why Prognostics and Health Management and Reliability, Availability, Maintainability and Safety have not married yet

Pierre Dersin

Operation and Maintenance Engineering Department, Luleå University of Technology, Sweden.
E-mail: pierre.dersin@ltu.se

Marielle Stoelinga

University of Twente and Radboud University, the Netherlands. E-mail: m.i.a.stoelinga@utwente.nl

RAMS (Reliability, Availability, Maintainability, Safety) and PHM (Prognostics and Health Management) are two engineering disciplines with different histories but broadly similar goals: managing risks resulting from failures, and mitigating them by acting on maintenance and operation.
Therefore the question naturally arises whether those two disciplines should be merged. This paper addresses that question by reviewing, for each of those two disciplines, its concept,its applications, its methods, and finally its benefits and limitations. Then we analyze their similarities and differences.

Keywords: RAMS, PHM, RUL, machine learning, model-based, data-driven, explainability, population, certification

1. Introduction

We argue that a tighter integration between Prognostics and Health Management (PHM) and Reliability, Availability, Maintainability, and Safety (RAMS) is instrumental to reap the benefits of predictive maintenance. RAMS and PHM are historically separated research fields although they share common purposes: preventing system failures by keeping assets in good health and optimizing maintenance. How those two disciplines approach failure prediction and uncertainty modeling differs, but those approaches are complementary. PHM algorithms focus on estimating and predicting the health of individual components, such as bearings or pumps, yielding customized component maintenance plans. Yet maintenance decisions must be made at the system level, which necessitates integrated strategies that account for interactions and dependencies among components. This is where RAMS methodologies are most effective since they effectively model component interactions and assess system-level risks. But traditional RAMS methodologies do not inherently support predictive maintenance. They traditionally rely on average, fleet level predefined operating conditions and fail to take advantage of real-time asset monitoring data. As a result, they can only model traditional maintenance strategies, typically scheduled preventive maintenance, and not dynamic predictive maintenance.

Thus, PHM and RAMS exactly complement each other: PHM provides insight in exact failure mechanisms of individual components and leverages monitoring data to predict actual component-level performance, while RAMS provides insight in components interaction and predicts average system-level performance. Despite that complementarity, evidenced in some recent publications (Moradi and Groth (2020)), as well as recent progress, integration of the two disciplines has not happened yet. The authors hope that the thoughts shared here will trigger useful discussions.

Organization of the paper. Section2 presents a brief survey of PHM advances, benefits and limitations. Section3 performs a similar survey for RAMS. Similarities and differences are described in Section 4. Finally, Section 5 offers a non-exhaustive list of ways to overcome challenges. Section 6 presents concluding remarks.

2. PHM

2.1. *PHM Concept*

The concepts of Prognostics and Health Management Gouriveau et al. (2016) have evolved and are now formalized in several standards, for instance IEEE std1856 822 (2017). The latter defines prognostics as “the process of predicting an asset’s remaining Useful Life (RUL) by predicting the rate of progression of a fault given the current assessment of degree of degradation, the load history, and the anticipated future operational and environment conditions, to estimate the time at which the asset will no longer perform its intended function within the desired specifications”. A key point is the reference to ‘intended function’ and ‘desired specifications”. Indeed, the concept of failure (IEC 60050-192) is synonymous with “loss of the ability to perform a function according to desired specification”. It is with that definition in mind that the reliability engineer says that an item of equipment has failed. For instance, a LED (light-emitting diode) has a function which is to generate light.Some physical parameters, mainly input current and temperature, will affect the efficiency of the LED, i.e., the amount of light generated by input current. When the efficiency drops under a certain threshold, the LED no longer operates according to specifications, and is therefore deemed failed. The threshold is determined by the user’s needs and may vary from one application to the next. For instance, in street lighting requirements may be less stringent than, say, for indoor lighting. This is valid only for progressive degradations, not sudden failures. In fact, PHM is suitable for items that are characterized by an increasing failure rate; sudden (unpredictable) failures correspond to a constant failure rate instead.The degradation state therefore contains more information than the “operational versus failed” state. A prerequisite, before performing a prognosis, is to detect an anomaly and to diagnose the particular degradation mode that is involved, since different degradation modes will generally evolve with different speeds to the failure. This is why PHM generally consists of three steps: anomaly detection; diagnostics; and prognostics. For instance, in LEDs, the main degradation mechanism is linked to quantum efficiency; degradation causes luminous flux reduction and color shift. With gearboxes (as found in trains or helicopters, for instance), degradation mechanisms include gear tooth wear. In train doors, degradation usually involve inefficiency of the kinematic chain (transmitting force from motor to actuators) due to wear. In turbofans of jet aircraft, degradations affect turbine efficiency and flow. Lithium-ion batteries also experience varopus physical aging mechanisms..

2.2. *PHM Applications*

With sufficient information on the evolution of degradations, one is armed to perform health management.In broad terms, health management consists of making decisions relating to maintenance and operations with the general goal of keeping assets operational as long as possible. Health management requires some knowledge of the future use of the system because that use will determine the stresses, and therefore future aging.Actions can include performing preventive maintenance, or more precisely condition-based maintenance or predictive maintenance, that is, maintenance triggered by the current condition of the asset or estimated *Remaining Useful life (RUL)*. Thus, predictive maintenance is an important application of PHM. It is however not the only one: reducing the stresses can lead to extending an asset’s life time; in fact, current research investigates the use of control algorithms Félix et al. (2023) to manage the RUL. Thereby PHM may impact not only maintenance, but also operations. The key input to PHM is health monitoring, through sensors that are either added specifically or pre-existing, as part of the control system for instance.Monitoring can be continuous but can also consist of regular inspections,and can rely on built-in test equipment as well.From the raw data acquired through monitoring,it is desired to estimate the health state of the monitored items, and, if possible, to predict its evolution. The set of data acquisition devices and data processing algorithms that enable that estimation and prediction is usually referred to as a PHM system.

2.3. *PHM methods*

PHM methods fall into three broad categories: model-based, data-driven and hybrid.

Model-based PHM methods. Purely model-based approaches—sometimes called "physics-of -failures", rely on mathematical representations of the physical degradation processes. A typical example is the Paris-Erdogan law for crack growth. Others include the Arrhenius model for temperature stress, Black's law (Black (2005)) which combines temperature and current stresses(such as for LEDs), or Norris-Landsberg for temperature cycling. Model-based approaches are limited by model accuracy and the imprecise knowledge of model parameters (e.g., activation energy).

Data-driven PHM methods. To overcome the limitations in model-based techniques, empirical data-driven approaches have gained popularity, especially with the tremendous progress in machine learning, and in particular deep learning, supported by high-performance hardware (Fink et al. (2020)). The most frequent algorithms rely on supervised learning, i.e., they are trained on data sets that correspond to normal operations and to a number of degradation modes which are labelled; and they are subsequently able, in principle, to detect abnormal operation (anomaly detection), to identify the particular failure or degradation mode (diagnosis); and,in the best cases,to perform predictions (prognosis). A typical tool of data-driven algorithms is "features", which are summaries of the acquired signals. For instance, statistical features such as variance, kurtosis, peak-to-peak distance (Atamuradov et al. (2020)), are commonly used. Feature fusion then can lead to the construction of health indices.The work of feature-engineering usually requires intensive expert involvement. Recent advances in machine learning, particularly the use of artificial neural networks, has enabled in some cases the replacement of feature engineering with feature learning, i.e. the algorithm learns suitable features automatically, and have made unsupervised learning possible, overcoming the need for labeling. Various types of auto-encoders (for instance variational auto-encoders) support that approach.

Hybrid PHM methods. Finally, hybrid methods combine data-driven and physics-based approaches. They include physics-informed machine learning (PIML), and in particular physics-informed neural networks (PINN) (Arias Chao et al. (2022)). Diverse combinations have been considered; from embedding physical equations into the neural network, to constraining the neural network by some high-level physics-based law (Bajarunas et al. (2023)).

2.4. *Benefits and Limitations of PHM*

Thus PHM is benefiting enormously from those recent very promising developments. In particular, complex nonlinear time-varying patterns can be investigated (as this is one of the hallmarks of artificial neural networks). At the same time, limitations must be recognized: 1) Data are often lacking (for instance for highly reliable items, which by definition fail rarely), or of insufficient quality. This will adversely impact the performance of the algorithms: false detections (false negatives and false positives), wrong diagnosis,inaccurate prognosis. 2) Intensive domain expert involvement is needed for model-based methods and supervised learning methods. 3) Explainability is often lacking with data-driven methods: the reasoning that has led the algorithms to its conclusion, so that the human maintainer or operator can be convinced, is not available (this area is currently a focus of research and good progress can be expected (see for instance Forest et al. (2024)) 4) Finally, as the approach is component-based, it usually does not consider the interactions between the components that make up a system. One last comment is that, while PHM algorithms focus on one item (say, an aircraft or an aircraft engine in a fleet of aircraft, a train or a train subsystem in a fleet of trains), health management decisions must consider the fleet implications.

Health monitoring tends to be costly. The business case is often challenging, unless the assets to be monitored are critical in terms of availability or maintenance costs. Then a convincing case can be made that PHM will bring about substantial, quantifiable improvements.

3. RAMS

3.1. *Concept*

RAMS stands for Reliability, Availability, Maintainability and Safety. It is an engineering discipline that offers methods and techniques for making systems more dependable. The field is closely related to Probabilistic Risk Assessment (PRA) and reliability engineering. These fields adopt a system-level perspective on the system's risks, exploiting engineering knowledge and logical system modeling, capturing how failures arise at component level and propagate through the system, causing system-level failures.

The main purpose of RAMS is to make risk-informed decisions regarding system design and operations. Typical design decisions concern the overall system architecture, the level of redundancy and the required quality of components. Operational decisions concern maintenance, such as inspection frequencies, spare parts management, repairs, and replacements, as well as health-aware control, such as stress reduction, and changes in operational profiles. This purpose can included other purposes, such as certification, documentation, and diagnosis.

RAMS methods are deployed in many different industries, such as nuclear power generation, rail, avionics and aerospace, automotive, water management. Moreover, several regulatory bodies require RAMS methods for certification.

Example 3.1. Figure 1 depicts a tiny example of a fault tree. For the top event (MeC; Medium Corrosion) to happen, two conditions are needed: the presence of water (WW), and a medium acid level (AcM). The latter can be achieved through the presence of either Hydrogen sulfide H_2S or oxygen O_2 or carbon oxide CO_2 is required.

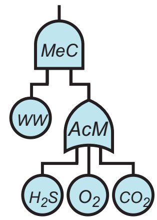

Medium Corrosion:	
Medium Corrosion	MeC
With Water	WW
Acid Medium	AcM
H_2S	H_2S
O_2	O_2
CO_2	CO_2

Fig. 1. Fault tree modeling corrosion in an oil-gas pipeline

3.2. *Applications*

Safety-critical domains include trains, planes, and nuclear power plants, where system reliability is paramount. Even in scenarios where safety is not the primary concern, availability often plays a crucial role—such as ensuring the smooth operation of services like passenger and goods transport, which heavily depend on consistent performance and uptime.

(1) **Understanding and documentation.** RAMS analysis provides a systematic overview of vulnerabilities and failure scenarios, enhancing understanding and serving as valuable documentation.
(2) **Compliance and certification.** RAMS analysis supports dependability metrics like reliability and MTTF, making it essential for demonstrating compliance in safety-critical systems.
(3) **Design and operational decisions.** By identifying cost-effective measures, RAMS analysis helps improve system dependability and reduce the probability of failures.
(4) **Diagnosis and monitoring.** RAMS analysis aids in tracing failure causes systematically, providing insights into how failures occurred.

3.3. *Methods*

Numerous RAMS methods exist. As in Boudali et al. (2007), we divide them into three classes.

Text-based methods. Textual approaches systematically explore components or behaviors in complex systems, presenting findings in textual form or tables. Common methods include failure mode effect analysis (FMEA) Rausand et al. (2020) and Hazard & Operability Studies (HAZOP) Kletz (1999). FMEA is often paired with fault tree analysis to identify components and failure modes as basic events.

Architectural methods. These methods use an architectural system model, i.e., a decomposition for the system into components and the way they interact. Such methods are common in systems with substantial software components, but apply to any complex design. In nuclear safety, the Figaro modeling language Bouissou et al. (1991) is a prominent example. Other examples include the AADL error annex Feiler et al. (2006), AltaRica Arnold et al. (1999), and Hip-HOPS Papadopoulos and McDermid (1999), often referred to collectively as Model-Based Safety Assessment Sun et al. (2024).

Domain-specific methods. These methods are tailored for specific risk analyses and include reliability block diagrams Modarres et al. (2009), event trees Ericson (2005), STAMP Leveson (2023), and bow tie diagrams Center for Chemical Process Safety (2018). Several dynamic variants of these models exist, enhancing their expressiveness and enabling the representation of common dependability patterns—such as spares and temporal orders—that cannot be captured by their original static counterparts.

3.4. *Analysis types*

The purposes mentioned in Section 4.2 are supported by two types of RAMS analyses:

(1) *Qualitative techniques* focus on identifying critical paths and failure causes. Common approaches include determining minimal cut sets and analyzing common cause factors.
(2) *Quantitative techniques* aim to calculate dependability metrics, which serve as key performance indicators for evaluating system dependability. A variety of analytical and statistical methods are available for this purpose. Typical metrics include: (a) The *System reliability*: the probability that a system operates without failure during its mission time. (b) *System availability*: the average proportion of time a system is operational. (c) *Mean time to failure (MTTF)*: the expected time until a system experiences its first failure.

3.5. *Benefits and Limitations of RAMS*

While RAMS excels at modeling component interactions and computing system-level risk metrics, it is not well-suited for predictive maintenance due to the following limitations

(1) **Lack of Monitoring Data Integration**: RAMS does not leverage monitoring data, instead focusing on average operating conditions or predefined scenarios. Moreover, automatic scenario synthesis based on evolving operating conditions or data analytics techniques is rarely addressed.
(2) **Limited Integration with Maintenance Practices**: RAMS typically relies on simplistic maintenance concepts, such as fixed repair times. As a consequence, maintenance decisions are based on average performance, while PHM approaches offer tailored just-in-time maintenance strategies to be effective.

4. PHM and RAMS: Similarities and Differences

Both RAMS and PHM share a common goal, which is the improvement of system availability and the cost-effectiveness of maintenance and operations. They both can rely to some extent on physical models. However, there are key differences, which are examined in more detail below.

(1) RAMS tends to model discrete events such as failures; while PHM models continuous processes such degradations leading to failures.
(2) Initially at least, CBM or PHM followed a “deterministic approach” e.g. developing indicators that trigger maintenance decisions, but without quantifying the corresponding uncertainty.
(3) RAMS is traditionally ”rational” and PHM empirical
(4) PHM focuses on individual assets, and RAMS on populations.

4.1. *Discrete versus continuous behaviour*

A key difference between PHM and RAM(S) is that, beyond the binary state “operational/failed”, PHM introduces a notion of degradation which progressively leads to a failure. This notion is illustrated in Figure 2: a health index (HI), a function of time, is defined, which measures the state of degradation (HI = 1 when there is no degradation, and HI= 0 corresponds to failure). Then the relationship with the reliability function is materialized by the relationship:

$$R(t) = P(HI(t) > 0] \tag{1}$$

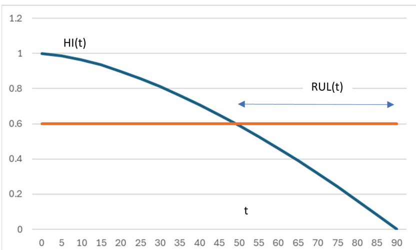

Fig. 2. Health Index and RUL as functions of time

4.2. *Applications*

However, beyond R(t), which is the concern of the reliability engineer, the PHM specialists are interested in HI(t) because they want to know how much time (or how many cycles or how many miles) is left before the failure threshold is hit; that is, the remaining useful life at time t, RUL(t). Knowledge of the RUL (or a good estimation of it) will support predictive maintenance policies.

4.3. *Deterministic versus Stochastic*

Reliability theory since its inception has been probabilistic by nature: among other, the very concept of reliability is defined by a probability. Field data statistics naturally include classical statistical uncertainty quantifiers such as confidence intervals. PHM has arisen from empirical condition monitoring and, until recently,the notion of statistical uncertainty was not included in its prediction, although this has recently changed.

4.4. *Rational versus empirical*

While RAMS typically relies on rational constructs (bottom-up such as FMEA or top-down such as FTA), PHM tends to follow empirical approaches, trying various algorithms and comparing their effectiveness on the same dataset. Explainability of the models is not necessarily inherent, especially with 'black-box' machine learning algorithms. This distinction applies to the system evaluation process, where the RAMS approach is mathematical while PHM rather uses empirical assessment methods.

4.5. *Asset versus Population*

The RAMS approach is population based, while PHM is customized to individual assets: RAMS considers populations of identical assets under supposedly identical operating conditions, as those assumptions are prerequisites for probabilistic and statistical approaches. For instance, MTTF (mean time to failure), MRL (mean residual life) or availability refer to a population, or an average asset in the population. In PHM, the relevant metrics are asset-specific RUL and HI. The MRL is the expectation of the RUL of all assets in a population. Thus, an intimate link exists between PHM and RAMS; the former acts at item level and the latter at population level (Dersin (2023)).

4.6. *Explainability versus Black box*

In many fields, safety demonstration and certification is a sine qua non condition for authorizing operations ; an important aspect of it is causal analysis: the causal chains that may lead to accidents must be understood in detail so mitigation measures can be taken and risks controlled.As seen earlier,empirical data-driven methods are often black boxes; there is now a strong drive toward explainability and causal inference but those are recent trends.

4.7. *Relation to Physics*

One area where RAMS and PHM approaches tend to coincide is their relation to physics. In RAMS, failure mode and effects analyses (FMEA) take into account the physical phenomena leading to degradations and failures; in PHM, purely data-driven approaches are increasingly complemented by physics-based models (Arias Chao et al. (2022)) as seen earlier. In reliability engineering, the physical degradation laws are used as well, through accelerated life tests Bagdonavicius and Nikulin (2019), but with a pre-defined mission profile which is deemed valid for an entire population. In PHM however, the mission profile is periodically or continually updated based on the sensor measurements on each specific item; therefore an estimate of the state of health and its evolution is available for a given specific asset, whose RUL is thus predicted. The notion of digital twin then embodies the bridge between the physical and the real world.

5. The integration of RAMS and PHM

Given the common goals, it would seem quite logical that both fields, RAMS and PHM, should merge. However, in industry at least, this has not happened so far. We believe the reasons to be partly historical and partly methodological: the two disciplines have evolved in different contexts and rely on people with different backgrounds, one community more mathematical and the other more empirical. The arrival of AI and ML, will probably contribute to accelerating the convergence, but so far ML has permeated PHM faster than RAMS; in particular, PHM application to safety-critical items encounters great obstacles due to insufficient explainability of the algorithms and rigidity of standards.

In reliability engineering the entire population is described by a time-to-failure (TTF) distribution, whose expectation is the MTTF; the conditional expectation of the remaining life of an asset of age t is the mean residual life (MRL) at time t. For the reasons already mentioned,including cost-benefit analysis, it seems therefore that a combination of RAMS-based fleet-level approach and PHM at individual item might be a promising avenue. For instance, RAMS-based estimates can provide prior distributions, in a Bayesian framework, for individual health indices.

6. Conclusion

The integration of PHM (Prognostics and Health Management) and RAMS (Reliability, Availability, Maintainability, and Safety) techniques shows significant potential but poses substantial challenges. This complexity arises from their fundamentally different methodologies: PHM emphasizes continuous degradation processes,while RAMS models failures as discrete events. PHM approaches can be deterministic (model-based) or statistical (data-driven), whereas RAMS techniques are inherently stochastic. Evaluation criteria also differ—PHM methods are assessed experimentally using confusion matrices and ROC curves, while RAMS methods are analyzed mathematically. Developing a unified framework necessitates harmonizing these contrasting perspectives. Recent trends show ways toward bridging the gaps. For instance the use of FMMEA (failure modes, *mechanisms* and effects analysis as the first step of the PHM process); the notion of *digital failure twin* Ge et al. (2023) ; inferring fault trees from data Jimenez-Roa et al. (2023); and Reliability-informed deep learning Dersin et al. (2024). In addition to technical challenges, there is also a cultural challenge which consists of breaking organization silos, particularly in industry.

Acknowledgements. The work reported in this paper was partially funded by the NWO under grants PrimaVera (NWA. 1160. 18. 238) and ZORRO (KICH1. ST02. 21. 003), Marie Skłodowska-Curie grant MISSION (101008233), ERC under grants CAESAR (864075) and RUBICON (101187945).

References

(2017). IEEE standard framework for prognostics and health management of electronic systems. *IEEE Std 1856-2017*, 1–31.

Arias Chao, M., C. Kulkarni, K. Goebel, and O. Fink (2022). Fusing physics-based and deep learning models for prognostics. *Reliability Engineering & System Safety 217*, 107961.

Arnold, A., G. Point, A. Griffault, and A. Rauzy (1999). The altarica formalism for describing concurrent systems. *Fundamenta Informaticae 40*, 109–124.

Atamuradov, V., K. Medjaher, F. Camci, N. Zerhouni, P. Dersin, and B. Lamoureux (2020). Machine health indicator construction framework for failure diagnostics and prognostics. *J.*

of signal processing systems 92, 591–609.

Bagdonavicius, V. and M. Nikulin (2019). *Accelerated Life Models-Modeling and Statistical Analysis*. CRC.

Bajarunas, K., M. Baptista, K. Goebel, and M. A. Chao (2023). Unsupervised physics-informed health indicator estimation for complex systems. In *Annual Conference of the PHM Society*, Volume 15.

Black, J. (2005). Mass transport of aluminum by momentum exchange with conducting electrons. In *2005 IEEE International Reliability Physics Symposium, 2005. Proceedings. 43rd Annual.*, pp. 1–6.

Boudali, H., B. R. Haverkort, M. Kuntz, and M. Stoelinga (2007). Best of three worlds: Towards sound architectural dependability models. In *8th International Workshop on Performability Modeling of Computer and Communication Systems (PMCCS)*, pp. 45–49.

Bouissou, M., H. Bouhadana, M. Bannelier, and N. Villatte (1991). Knowledge modeling and reliability processing: presentation of the Figaro language and associated tools. In *Proc. of Symposium on Safety of Computer Control Systems (SAFECOMP)*, pp. 69–75.

Center for Chemical Process Safety (2018). *Bow Ties in Risk Management*. Wiley.

Dersin, P. (2023). *Modeling Remaining Useful Life Dynamics in Reliability Engineering (1st ed.)*. Taylor and Francis/CRC.

Dersin, P., K. Bajarunas, and M. Arias-Chao (2024). Analytical modelling of health indices for prognostics and health management. In *Vol. 8 No. 1 (2024): Proceedings of the European Conference of the PHM Society*.

Ericson, C. A. (2005). *Event Tree Analysis*, Chapter 12, pp. 223–234. WILEY.

Feiler, P. H., D. P. Gluch, and J. Hudak (2006). The architecture analysis & design language (AADL): An introduction. Technical report, Carnegie Mellon University.

Fink, O., Q. Wang, M. Svensén, P. Dersin, W.-J. Lee, and M. Ducoffe (2020). Potential, challenges and future directions for deep learning in prognostics and health management applications. *Eng. Appl. Artif. Intell. 92*, 103678.

Forest, F., H. Porta, D. Tuia, and O. Fink (2024). From classification to segmentation with explainable ai: A study on crack detection and growth monitoring. *Automation in Construction 165*, 105497.

Félix, M. S., J. J. Martinez, and C. Bérenguer (2023). A state-space approach for remaining useful life control. *IFAC-PapersOnLine 56*(2), 7728–7733. 22nd IFAC World Congress.

Ge, X., Z. Zeng, and N. Anwer (2023). A digital failure twin model for phm: From concepts to maturity levels. *2023 Global Reliability and Prognostics and Health Management Conference (PHM)*, 1–6.

Gouriveau, R., K. Medjaher, and N. Zerhouni (2016). *From prognostics and health systems management to predictive maintenance 1: Monitoring and prognostics*. Wiley.

Jimenez-Roa, L. A., T. Heskes, T. Tinga, and M. Stoelinga (2023). Automatic inference of fault tree models via multi-objective evolutionary algorithms. *IEEE Trans. Dependable Secur. Comput. 20*.

Kletz, T. A. (1999). *Hazop & Hazan: Identifying and Assessing Process Industry Hazards* (fourth ed.). CRC Press.

Leveson, N. (2023). *An Introduction to System Safety Engineering*. MIT Press.

Modarres, M., M. Kaminskiy, and V. Krivtsov (2009). *Reliability Engineering and Risk Analysis: A Practical Guide*. CRC Press.

Moradi, R. and K. M. Groth (2020). Modernizing risk assessment: A systematic integration of pra and phm techniques. *Reliability Engineering and System Safety 204*, 107194.

Papadopoulos, Y. and J. A. McDermid (1999). Hierarchically performed hazard origin and propagation studies. In *Computer Safety, Reliability and Security*, pp. 139–152. Springer.

Rausand, M., A. Barros, and A. Hoylan (2020). *Qualitative System Reliability Analysis*, Chapter 4. Wiley.

Sun, M., S. Gautham, Q. Ge, C. Elks, and C. Fleming (2024). Defining and characterizing model-based safety assessment: A review. *Safety Science 172*.