How to Assess the Technology Readiness Levels of an AI-based System?

Stéphanie Gourdin
Thales Defense Mission Systems,
Elancourt, France
stephanie.gourdin@fr.thalesgroup.com

Pierrick Richard
Thales Defense Mission Systems,
Elancourt, France
pierrick.richard@fr.thalesgroup.com

Juliette Mattioli
Thales SA, cortAIx
Palaiseau, France
juliette.mattioli@thalesgroup.com

***Abstract*—The Technology Readiness Level (TRL) framework provides a rigorous yet adaptable structure for evaluating the maturity of Artificial Intelligence (AI) based systems, ensuring that advancements progress from theoretical research to operational deployment in a measured and transparent manner. Unlike traditional technologies, AI systems demand a more nuanced assessment due to their reliance on dataset and knowledge-base quality, AI model adaptability, and dynamic performance under varying conditions. In the early research phase (TRL 1–3), the focus lies in defining foundational elements such as operational boundaries, data governance, and preliminary compliance with regulatory standards like the AI Act, thereby mitigating future risks of non-compliance or technical shortcomings. As development advances into TRL 4–6, validation extends beyond laboratory settings to real-world environments, where AI components must demonstrate not only functional correctness but also robustness, scalability, and hardware compatibility. The final deployment phase (TRL 7–9) emphasizes full-system integration, ethical alignment, and sustained operational reliability, ensuring that AI solutions meet both technical and regulatory benchmarks before widespread adoption. In this paper we discuss the adjustment needed in the TRL evaluation for AI based systems and share a structured checklist-based tool to support this process by defining criteria and advisory risk-mitigation measures, fostering a balanced approach to innovation while preventing costly oversights.**

Keywords-* *AI system; Technology Readiness Level; TRL criteria; AI maturity assessment

I. INTRODUCTION

High-quality software products and computer systems are crucial to stakeholders. Quality models, quality requirements, quality measurement, and quality evaluation are standardized within the International Standards on SQuaRE [1]. The Technology Readiness Level (TRL) framework [2][3] offers a systematic approach for assessing the maturity of technological innovations, providing a common language that bridges the gap between technical development and strategic decision-making. By establishing clear, measurable criteria for assessing progress, the TRL scale enables organizations to communicate effectively across disciplines, ensuring that engineers, project managers, and executives share a unified understanding of where a technology stands in its development lifecycle. This shared framework not only facilitates more informed discussions about technical maturity, but also enhances the ability to conduct comprehensive risk assessments, allowing stakeholders to identify potential challenges and dependencies before they become critical issues [4]. Furthermore, the TRL approach serves as a valuable guide for investment decisions, helping organizations allocate resources more efficiently by distinguishing between technologies that require additional development and those that are already mature enough for deployment. In this way, the framework prevents the premature introduction of unproven solutions while simultaneously avoiding excessive investment in technologies that have already reached their full potential.

In order to guarantee the various trust properties of an AI-based system throughout its lifecycle [5][6][7], it is important to assess its TRL. Unlike conventional deterministic software, AI systems exhibit emergent behaviors derived from data and knowledge rather than explicit programming. Without rigorous validation, their performance characteristics and potential failure modes cannot be fully understood. Given the rapid evolution and unique challenges of AI development, it is crucial to adapt and refine maturity assessment methods such as the Technology Readiness Level (TRL) framework to ensure the responsible and effective deployment of AI technologies.

Today, some standards define maturity criteria [8][9], but to our knowledge there is no practical guideline on how to carry out such an assessment in practice. In this article, we present a dedicated tool developed to facilitate the adaptation and application of the TRL framework for Artificial Intelligence (AI) systems. The first section provides a brief overview of the TRL scale and discusses the reasons why adjustments are necessary when applying it to AI systems. Next, we describe the methodology used to derive the proposed framework and demonstrate its use in practice through our tool. We then detail the AI-specific adjustments made at each main stage of the TRL framework—namely, the research phase, development phase, and deployment phase—before concluding.

II. TECHNOLOGY READINESS LEVEL FRAMEWORK IN THE CONTEXT OF ARTIFICIAL INTELLIGENCE

The nine-level TRL scale provides a common language for evaluating and communicating the development status of technologies, from initial concept through full operational deployment [10]. The scale progresses through three broad phases:

- Research Phase (TRL 1-3): Basic principles are observed and reported (TRL 1), technology concepts are formulated (TRL 2), and proof of concept is established through analytical or experimental means (TRL 3). This phase focuses on fundamental research and feasibility studies.
- Development Phase (TRL 4-6): Component validation occurs in laboratory environments (TRL 4), followed by validation in relevant environments (TRL 5), and

culminating in system demonstration in relevant environments (TRL 6). This phase involves prototyping, testing, and integration of components into increasingly realistic conditions.

- Deployment Phase (TRL 7-9): System prototypes are demonstrated in operational environments (TRL 7), actual systems are completed and qualified through testing and demonstration (TRL 8), and finally, systems are proven through successful mission operations (TRL 9). This final phase represents the transition from development to full operational capability.

The framework provides a clear roadmap for technology progression, facilitating risk management, resource allocation, and alignment with technical and strategic objectives. This is particularly relevant in sectors where safety and reliability are critical, such as aerospace, defense, cybersecurity, and digital technology. For Artificial Intelligence (AI) systems [11], a TRL assessment is not just a formality, but a key factor in determining feasibility, risk and investment decisions.

The EU AI Act defines "*AI systems as machine-based systems operating with varying autonomy, adapting over time, and generating influential outputs, from predictions to decisions*". The field of AI mainly follows three major paradigms. The first is *data-driven AI*, which covers statistical and connexionist AI such as Machine Learning (ML) and Generative AI. Inspired by biological neural networks, this sub-discipline has dominated since the early 2020s due to its ability to deduce patterns from data. Secondly, *knowledge-based AI* (also known as symbolic AI) relies on knowledge representations such as ontologies and conceptual graphs. The distinction between the first two paradigms lies in how knowledge is acquired: data-driven AI extracts it automatically from examples, whereas symbolic AI encodes it explicitly through human expertise. Finally, *hybrid AI* encompasses any synergistic combination of various AI techniques, which could be enhanced by prior knowledge (such as mathematics, physics, or geometry).

Unlike traditional engineering systems, where TRLs are often tied to physical prototypes and testing, AI systems introduce unique complexities due to their adaptive, data-driven and frequently opaque nature. For example, AI-based systems may demonstrate exceptional performance within the controlled parameters of a laboratory setting, achieving metrics that suggest a high level of readiness [12][13]. However, when exposed to the unpredictability of real-world conditions, such as adverse weather, adversarial threats or the variability of human behavior, their effectiveness can diminish considerably. AI-specific failure modes must be carefully addressed before AI algorithms are deployed. For instance, data-driven AI models can become miscalibrated due to subtle shifts in data distribution during deployment, causing them to overestimate their predictive accuracy. The disparity between controlled testing and real-world application underscores the pivotal role of performance and TRL assessments. These provide a more accurate understanding of a technology's readiness.

The disparity between AI performance in controlled lab settings and real-world applications highlights the need to manage stakeholder expectations effectively. These expectations are frequently influenced by overly optimistic forecasts or marketing claims rather than concrete evidence. TRL assessments [14] provide an unbiased evaluation of an AI system's capabilities, ensuring that stakeholder expectations align with actual performance. This alignment fosters greater transparency and builds trust in the technology's potential.

TRL evaluations also guide Research and Development (R&D) by identifying the specific challenges that hinder an AI system's transition from theoretical success to practical implementation. For instance, a medical diagnostic AI tool may demonstrate high accuracy when analyzing historical patient data, suggesting advanced readiness in a controlled environment. However, when deployed in real clinical settings, it may encounter unexpected obstacles, such as compatibility issues with hospital infrastructure or usability limitations, that reduce its operational effectiveness. By pinpointing these gaps, TRL assessments enable organizations to prioritize efforts toward overcoming critical barriers, whether through improving system interoperability, enhancing user experience, or refining the technology's scope to more feasible applications.

Furthermore, the regulatory and ethical implications of AI deployment underscore the importance of rigorous TRL evaluations, particularly in high-stakes sectors where safety and security are paramount. Applications like autonomous weapons or AI-driven medical diagnostics demand thorough validation to ensure compliance with ethical and regulatory standards. For example, an AI-based credit scoring system must not only prove technical proficiency but also meet strict fairness, robustness, and transparency requirements before being responsibly integrated into financial institutions. The TRL framework offers a structured method for assessing these multifaceted dimensions, ensuring AI systems achieve the necessary certifications and public trust. This approach helps organizations navigate complex regulatory landscapes while mitigating risks associated with premature or inadequately validated deployments.

III. METHODOLOGY FOR DEVELOPING TRL EVALUATION CRITERIA FOR AI

The goal of this work is to establish clear criteria to objectively determine each TRL level, as the interpretation of the TRL scale can vary. While specific criteria could also be useful for evaluating classical algorithms, experts have generally relied on their experience and a shared understanding to assess maturity without a detailed scale. In contrast, the use of new AI-based algorithms—whether data-driven or knowledge-based— necessitates a reevaluation of how algorithm maturity is assessed and calls for more precise and tailored criteria.

Indeed, since these systems are less reproducible and data-dependent, there is a greater reliance on testing methodology and environment: applying classical criteria alone may lead to an overestimation of the maturity level. For example, using real data for both training and testing in volumes and with representativeness that are insufficient to ensure proper algorithm generalization can give misleading results. Some AI-based algorithms can produce convincing results very quickly, but their maturity also depends on the maturity of the testing environment. Therefore, we propose to precisely detail the TRL criteria to ensure the actual maturity of algorithms. Here, maturity is defined as the ability to use the algorithm within a system without risk of unexpected or untested behavior.

The methodology to develop the AI TRL scorecard was based on various sources proposing approaches for measuring TRLs in classical algorithms, used by different institutions to define the meaning of the different levels on the scale (including documents from French Department of Defense, NATO, the European Union, and the ISO 16290 standard [8]). From this work, it emerged a list of criteria that can be used to measure TRL levels for classical algorithms more precisely than by merely interpreting the level definitions alone.

We based our work on Thales' dual experience: on one hand, in our algorithmic studies aimed at integrating Artificial Intelligence into our products within the demanding defense context; on the other hand, through our participation in the Confiance.ai program [15], which highlights critical issues that may arise during the maturation of algorithms whose definitions are not derived from explainable problem modeling but rather from learning. This work has allowed us to adapt and supplement the list of criteria to be considered specifically for the challenges posed by Artificial Intelligence.

The tool helps project teams assess algorithm maturity and identify early risks. Using a checklist in an Excel template, it shows the current TRL level and highlights areas needing improvement. It also serves as a record to document progress and support maturity claims.

Mandatory criteria directly contribute to TRL definitions by characterizing algorithm maturity; failure to meet any mandatory criterion indicates the associated TRL level is not achieved. These criteria align with ISO 16290 requirements for TRL validation.

Conversely, recommended criteria function as risk mitigation checkpoints to ensure critical considerations are addressed timely, preventing potential deadlocks during further maturation phases that could necessitate solution redesign. While strongly advised to be reviewed at specified milestones, non-compliance with these does not invalidate the reported TRL. For instance it is recommended to pay attention to software licensing: during the research phase, the license used does not impact the progression in maturity, but it may prevent industrialization if incompatible. Similarly, regulatory compliance (e.g., AI Act) does not impede research but prohibits product deployment if breached. Therefore, it is advisable to address compliance issues as early as possible; otherwise, investments in the technology may turn out to be futile, as the solution could turn onto a non-deployable technology.

IV. Methodology for Assessing AI TRL

To assess the Technology Readiness Level (TRL), it is essential to clearly define and baseline the technology or subsystem being evaluated. This involves specifying what is being assessed, its intended use, and its boundaries—a definition that becomes more detailed as the system matures. For higher TRLs, it is also necessary to clearly state the performance requirements and understand the mission, the system context, and the operational environment involved [8].

The evaluation tool includes a dedicated section for describing the element under assessment, with the level of detail adapted to the maturity phase (research, development, or deployment). For Artificial Intelligence projects, this means specifying both the algorithm's objectives and the details of its training process, especially the datasets used. Since AI development is highly iterative, with frequent changes even late in the process, TRL assessment for AI must also be dynamic and ongoing. Regular updates to the training database are common to improve the system's robustness and alignment with real-world conditions.

A TRL score alone does not provide an exact measure of the remaining effort or costs, especially for AI, where initial development (such as dataset creation and labeling) can be resource-intensive, but subsequent updates may be efficiently revalidated through automated testing frameworks.

Achieving TRL 9 requires reproducible algorithm behavior, which is often challenging for AI systems that naturally incorporate randomness or adaptivity. Small updates usually only impact specific criteria, and a detailed evaluation helps identify what must be retested versus what remains unaffected by changes. Adopting an iterative and flexible approach is key: accepting temporary regressions in TRL fosters continuous improvement and shorter validation cycles. Instead of focusing solely on linear TRL progression, this mindset supports faster refinements and helps accelerate industrial maturity, balancing readiness assessment with ongoing enhancement.

V. Research Phase (TRL 1-3)

In the classical interpretation of TRLs, the early research phase (TRL 1–3) focuses on identifying fundamental principles, defining an application concept, and demonstrating technical feasibility through a proof of concept. For conventional software or hardware systems, readiness at these levels is primarily assessed through the identification of system inputs and outputs, the definition of functional requirements, and the verification that the proposed solution complies with applicable standards and regulations. However, when dealing with AI-based systems, these criteria must be significantly extended to account for the data-driven and probabilistic nature of learning-based components or the stochastic nature of certain knowledge-based AI or hybrid AI approaches.

At TRL 1, beyond simply identifying system inputs and outputs, it becomes necessary to characterize the nature of the data and knowledge itself [16]. We recommend assessing early on, starting at this initial stage, whether the underlying AI technics principles comply with AI-specific regulatory constraints, such as those outlined in the AI Act [17]. This early evaluation helps avoid investing effort in developing systems that may later prove non-compliant.

% Completed	44 %		TRL 3	Analytical and/or experimental demonstration of the feasibility of critical functions
—	100	☑	Mandatory	At least one feasible application is identified for the specified technological concept and associated basic principles.
—	100	☑	Mandatory	The functional chain of the technology is established
—	0	☐	Mandatory	If the algorithm is data-driven, the data lifecycle is established: creation (data + labels), storage, data quality, and accessibility. If it is knowledge-driven, the knowledge lifecycle is established
—	100	☑	Mandatory	The evaluation database covers the domain of application of the algorithm in its core functions.
—	0	☐	Mandatory	Simulations and/or laboratory experiments on subsets have justified the feasibility of critical elements of the technology for the intended application.
—	100	☑	Mandatory	Performance measurement metrics are established, covering both AI-specific metrics and operational metrics of the target application. Metrics ensuring test statistics are included.
—	0	☐	Mandatory	These simulations have confirmed/clarified/optimized the expected performance magnitudes.
—	0	☐	Mandatory	The lifecycle of AI algorithm design is mastered, and the ability to reproduce performance is achieved.
—	0	☐	Mandatory	The algorithm's complexity is calculated with a view to implementation goals compatible with target constraints.
		☐	Recommended	Performances gains over previous and competing technologies are confirmed.
		☐	Recommended	Major risks and blocking points are identified (technology integration, environments, etc.).
		☐	Recommended	A document justifying choices or a study report is documented. (A document justifying choices or a study report is documented.)
		☐	Recommended	A potential user (client, product line) has expressed a preliminary interest in one of the identified applications within the scope and level of generality decided. This user will subsequently be referred to as the "client."
		☐	Recommended	The metrics for algorithm explainability are identified.
		☐	Recommended	The metrics for robustness and/or the availability of a monitoring tool for the algorithm are identified.
		☐	Recommended	The skills and processes (methods, tools, manufacturing, etc.) necessary for the development of the technology are identified.
		☐	Recommended	A roadmap towards concrete applications is outlined.

Figure 1. Criteria for TRL3

TRL 2, traditionally focused on defining the application domain and conceptual system architecture, must explicitly introduce the Operational Design Domain (ODD) [18] and formalize the processes for obtaining, qualifying, and governing data and knowledge, since their availability and representativeness directly impact system feasibility. This stage is achieved through initial simulations aimed at demonstrating the concept, without yet addressing representativity.

At TRL 3, where a proof of concept is established, readiness for AI-based systems cannot be inferred solely from functional demonstrations. Instead, it requires explicit evidence that evaluation criteria tailored to AI-based behavior have been formally defined, including performance metrics that capture AI system effectiveness, uncertainty, and failure modes, and whose computation is reproducible across datasets and experimental runs [12]. These metrics must be supported by a structured data lifecycle encompassing data ingestion, curation, and labeling processes, and complemented by initial considerations regarding explainability and robustness, which are necessary to interpret and contextualize measured performance.

At the end of this research phase, the explicit label of the technology, the level of generality and the research and industry players developing the technology is fixed and are part of the context for the evaluation. Furthermore, the functional chain of the technology and the performance measurement metrics are also established and baselined.

Consequently, for AI-based systems, TRL 1–3 no longer represent a simple progression from concept to feasibility, but rather a structured reduction of uncertainty across data, learning behavior, and evaluation criteria.

We show in Figure 1 how the template is constructed for the example of the TRL3 which is a major step especially for AI systems.The example particularly illustrates a case where some of the criteria required to validate the TRL are met, but not all. In this case, the template allows measuring the level of TRL achievement as a percentage. The criteria shown in black are common to both classical and AI algorithms, while the criteria specific to Artificial Intelligence work are highlighted in orange.

VI. Development Phase (TRL 4-6)

In the classical TRL framework, the development phase spanning TRL 4 to TRL 6 corresponds to the progressive validation of system components and their integration under increasingly realistic conditions. TRL 4 focuses on the validation of individual components or subsystems within a laboratory environment, TRL 5 extends this validation to relevant environments that approximate operational conditions, and TRL 6 culminates in the demonstration of a system or subsystem prototype operating in a representative environment. For conventional systems, readiness at these levels is assessed primarily through functional correctness, interface compatibility, and the ability to integrate components into a coherent system architecture.

For AI-based systems, however, this development phase introduces additional and non-trivial maturity criteria that go beyond classical component validation. At TRL 4, validation is not limited to functional behavior in a laboratory setting: the AI architecture must be explicitly designed to be compatible with the target hardware constraints, such as computational

 ISBN: 978-1-68558-395-8

% Complete	0 %	TRL 6	Demonstration using a demonstrator or prototype in a highly representative environment
—	0 ☐	Mandatory	The operational environment of the system is known.
—	0 ☐	Mandatory	A set of technology requirements is developed with the client.
—	0 ☐	Mandatory	A prototype that meets all these requirements is created.
—	0 ☐	Mandatory	A satisfactory demonstration of the prototype in the main non-laboratory environmental conditions required by the client has been performed.
—	0 ☐	Mandatory	The client and the supplier agree on the representativeness of the client's demonstration needs.
—	0 ☐	Mandatory	The demonstration has validated the architecture, functions, integration of risk aspects, and system specifications and its subsystems.
—	0 ☐	Mandatory	An analysis of real-time constraints is conducted.
	☐	Recommendec	Specific requirements regarding robustness against cyber attacks (e.g., adversarial) are identified and included in the tests.
	☐	Recommendec	Procurement and manufacturing processes have been tested and validated through the prototype's experience feedback.
	☐	Recommendec	All technology elements, including its interfaces, are available or easy to develop without difficulty.

Figure 2. Criteria for TRL6

resources, memory, and latency requirements. Moreover, the AI architecture and its interfaces are formally specified, enabling controlled integration with non-AI components. A critical addition at this level is the creation and use of a dedicated validation or qualification dataset, distinct from training datasets or the knowledge base, to assess AI model behavior under controlled and repeatable conditions.

At TRL 5, the transition to a relevant environment requires a stronger stabilization of the AI component itself. In contrast to earlier stages where iterative design of the AI algorithm design remains central, the model is considered ready for deployment in the sense that its parameters are frozen, thereby fixing the AI capacity behavior. The evaluation dataset is explicitly determined based on client or stakeholder needs, reflecting operational expectations and acceptance criteria rather than research-driven objectives. Performance validation at this stage relies on simulation environments or real-world data that were not used during training, in order to assess generalization capabilities and to reduce the risk of overly optimistic performance estimates.

The TRL 6, detailed in Figure 2, further extends classical system demonstration by introducing security considerations that are specific to AI-based systems. In addition to demonstrating functional performance in a representative environment, readiness at this level requires that the system's robustness against adversarial threats is explicitly taken into account. This includes assessing the susceptibility of the AI model to adversarial inputs or malicious perturbations and evaluating the impact of such attacks on system behavior [19]. Consequently, for AI-based systems, TRL 4–6 represent not only a progression toward system integration, but also a shift from AI-centric development to controlled deployment readiness, where architectural compatibility, dataset governance, model stability, and security become central indicators of technological maturity.

VII. Deployment Phase (TRL 7-9)

In the deployment phase, the objective is to transition from a prototype to a product that can be industrialized. The algorithms have already demonstrated their performance and robustness during the previous maturity ramp-up phases, and therefore no further modifications are made to them starting from this deployment stage. At this point, the Artificial Intelligence algorithms are fixed, and the process of moving from prototype to industrialization are identical to those applied to traditional software projects and do not require AI-specific adjustments.

VIII. Conclusion and Future Works

This paper presents a TRL assessment guide which offers a structured framework in order to guarantee the various trust properties of an AI-based system throughout its lifecycle. As operational contexts evolve, data-driven AI systems can degrade in ways traditional technologies do not. TRL assessment therefore verifies that systems are validated under sufficiently representative conditions, preventing premature advancement based only on selected benchmarks.

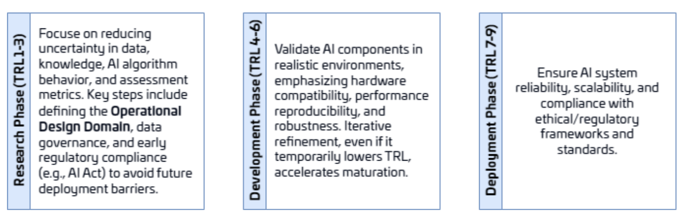

Figure 3. The TRL framework supports the AI system maturity assessment, bridging technical development and strategic decision-making.

As resumed Figure 3, lower TRLs address core capabilities, intermediate TRLs show subsystem integration, and higher TRLs validate fully integrated systems. AI components must also fit into complex existing architectures. The stepwise TRL progression ensures integration issues are found at the right development phase, not at deployment.

The TRL framework provides a common language for AI developers, users, acquisition authorities, and decision-makers. It helps prevent misalignment, such as laboratory demonstrations creating unrealistic expectations of operational maturity, and is essential for governance and multinational collaboration.

This TRL framework is already used at Thales in various operational and research contexts [20]. It supports defining functional requirements and evaluation metrics during proposals, assessing the maturity of research components before transfer to Thales solutions, and ensuring that product and system maturity suits their intended defense and aeronautics uses. As noted, it is integrated into our engineering workflows to systematically reinforce AI governance. The tool is non specific to Thales and easy to use, therefore in the mid term, it will be made Open Source as part of the European Trustworthy AI Association[1], which provides open-source tools to support scalable and secure AI development.

The TRL scale nonetheless has limitations. Although it details the path from research to large-scale deployment, it does not inherently cover other factors critical to success. Regulatory compliance and market maturity add complexity: technologies must both function correctly and meet legal standards while being accepted in target markets. These issues are acute when moving between sectors, where differing interpretations and adaptations of TRLs can cause inconsistencies and misunderstandings.

In safety-critical fields such as aerospace and defense, where AI failures can be catastrophic, adherence to the TRL process is both best practice and an ethical imperative [21]. Each TRL stage functions as a checkpoint, requiring evidence of increased technological maturity before progression. This staged, evidence-based approach ensures rigorous validation under realistic conditions before moving from laboratory prototypes to operational use. It is especially crucial for systems that must satisfy certification standards like EUROCAE/SAE ARP6983 for aeronautics [22]. Emerging regulations, such as the European AI Act, also mandate phased development with progressive demonstrations of maturity, for which TRL assessments provide a natural structure, particularly in high-risk applications.

References

[1] ISO/IEC 25000:2014, *Systems and software engineering — Systems and software Quality Requirements and Evaluation (SQuaRE) — Guide to SQuaRE*, Standard, 2014.

[2] J. Mankins et al., *Technology readiness levels*, 1995.

[3] S. Hirshorn et al., *Nasa systems engineering handbook*, 2017.

[4] S. Li et al., "A multi-dimensional assessment system for technology readiness levels", in *2017 4th International Conference on Systems and Informatics (ICSAI)*, IEEE, 2017, pp. 798–802.

[5] A. Awadid et al., "AI systems trustworthiness assessment: State of the art", in *12th International Conference on Model-Based Software and Systems Engineering*, 2024.

[6] K. Quintero et al., "An end-to-end method for operationalizing trustwothiness in AI-based critical systems", in *15th International Conference on Performance, Safety and Robustness in Complex Systems and Applications PESARO 2025*, 2025.

[7] E. Amdouni et al., "An end-to-end trustworthy knowledge graph engineering methodology", in *16th International Conference on Performance, Safety and Robustness in Complex Systems and Applications PESARO 2026*, 2026.

[8] *ISO 16290 : Space systems - Definition of the Technology Readiness Levels (TRLs) and their criteria of assessment*, Standard, 2014.

[9] B. Tucker and M. Bailey, "Rethinking technological readiness in the era of ai uncertainty", *arXiv e-prints*, arXiv–2506, 2025.

[10] J. Mankins, "Technology readiness assessments: A retrospective", *Acta Astronautica*, vol. 65, no. 9-10, pp. 1216–1223, 2009.

[11] J. Mattioli and C. Meyer, "Artificial Intelligence for Critical Systems", Thales, Tech. Rep., 2024.

[12] J. Mattioli et al., "An overview of key trustworthiness attributes and KPIs for trusted ML-based systems engineering", *AI and Ethics*, vol. 4, no. 1, pp. 15–25, 2024.

[13] J. Mattioli et al., "A brief overview of key quality metrics for knowledge graph solution illustration on digital notams", in *Proceedings of the AAAI Symposium Series*, vol. 7, 2025, pp. 206–213.

[14] A. Lavin et al., "Technology readiness levels for machine learning systems", *Nature Communications*, vol. 13, no. 1, p. 6039, 2022.

[15] Confiance.ai et al., *Towards the engineering of trustworthy AI applications for critical systems - the confiance.ai program*, 2022.

[16] J. Mattioli et al., "Information Quality: the cornerstone for AI-based Industry 4.0", *Procedia Computer Science*, vol. 201, 2022.

[17] L. Floridi, "The European legislation on AI: a brief analysis of its philosophical approach", *Philosophy & Technology*, vol. 34, no. 2, pp. 215–222, 2021.

[18] P. Koopman and F. Fratrik, "How many operational design domains, objects, and events?", in *Safeai@ AAAI*, 2019.

[19] M. Gonzalez et al., "Introducing RUM: A Methodological Contribution for Engineering Trustworthy AI Components in Industrial Systems", in *Proceedings of the AAAI Symposium Series*, vol. 7, 2025, pp. 153–160.

[20] *Thales AI in Real Use*, 2025. [Online]. Available: https://www.thalesgroup.com/en/advanced-technologies/artificial-intelligence.

[21] M. Cummings, "Rethinking the maturity of artificial intelligence in safety-critical settings", *AI Magazine*, vol. 42, no. 1, pp. 6–15, 2021.

[22] *ARP6983 - Process Standard for Development and Certification/Approval of Aeronautical Safety-Related Products Implementing AI*. [Online]. Available: https://www.sae.org/standards/arp6983-process-standard-development-certification-approval-aeronautical-safety-related-products-implementing-ai.

[1] https://www.trustworthy-ai-association.eu/