Carnegie Mellon University
Software Engineering Institute

A Preliminary Report on a Model for Maturing AI Adoption: From Hype to Achieving Repeatable, Predictable Outcomes

Ipek Ozkaya
Anita Carleton
Matt Butkovic
Sebastian Echeverria
Robert Edman
John Haller
Erin Harper
Mike Konrad
Natalie Schieber
Carol Smith
Shawn Wray

WHITE PAPER

DOI: 10.1184/R1/30840476

December 2025

DM25-1535

Table of Contents

Executive Summary **ii**

1 Introduction **1**
- 1.1 Hype Versus Reality 1
- 1.2 Why Organizations Fail at AI Adoption 2
- 1.3 Driving Maturity Through Predictable, Repeatable Measurement and Analysis 3
- 1.4 AI Is Software 4

2 State of the Practice in AI Maturity **6**
- 2.1 Review of AI Maturity Modeling Efforts 6
- 2.2 Building on a History of Maturity Modeling 7

3 The AI Adoption Maturity Model **8**
- 3.1 Building Blocks of the AI Adoption Maturity Model 9
- 3.2 Eight Core Dimensions of Capability 9
  - 3.2.1 Organizational Change Dimensions 9
  - 3.2.2 Lifecycle Engineering Dimensions 10
- 3.3 An Adaptable Maturity Assessment Framework 11
- 3.4 Progressing in Maturity: AI Adoption Maturity Model Levels 13
  - 3.4.1 Exploratory AI 14
  - 3.4.2 Implemented AI 15
  - 3.4.3 Aligned AI 16
  - 3.4.4 Scaled AI 17
  - 3.4.5 Future Ready AI 18
- 3.5 Roadmap Development 19

4 Conclusion and Next Steps **20**

Bibliography **21**

Executive Summary

Fast-evolving advances in artificial intelligence (AI), particularly in generative AI technologies and agentic AI implementations, have challenged organizations to rethink how they conduct business. There is an increasing frenzy to take advantage of AI not only in the hope of transformational productivity and cost improvements, but also to increase the pace of innovation. However, many organizations are failing to see the promised positive outcomes for their AI efforts, especially as they incorporate generative AI into their workflows and products. Lessons learned reveal common challenges, including a combination of mismatched expectations, misaligned applications, and poorly executed or untested implementation practices. This points to a common gap: industry and the Department of War both lack an adaptable, measurement-based tool for consistently assessing the maturity of their capabilities, enabling predictable AI adoption, and building a roadmap to strengthen practices needed to achieve their AI goals.

To meet this need, the Carnegie Mellon University Software Engineering Institute (CMU SEI), in partnership with Accenture, is developing the AI Adoption Maturity Model. This document introduces the key concepts of the model, which will provide organizational leaders with guidance on overcoming the challenges that arise as they try to realize the promise of AI. The model allows organizations to measure the degree to which key practices are implemented and governed for adapting and delivering AI solutions that meet business needs with predictable outcomes. The goal of the model is to drive the creation of a roadmap based on an assessment for successful AI adoption. The AI Adoption Maturity Model is being created with a thorough research and development effort, including executive interviews, a systematic review of over a hundred existing AI maturity efforts, pilots of AI projects, an ongoing survey to collect industry experiences, the SEI's extensive expertise in maturity modeling, and Accenture's hands-on experience worldwide with AI implementation.

This report describes the proposed model's core concepts and five levels of AI adoption maturity and their specific characteristics. These include maturation through the levels of Exploratory AI, Implemented AI, Aligned AI, Scaled AI, and Future Ready AI. This report details the eight core organizational change and system lifecycle engineering dimensions and exemplifies their associated capabilities that organizations need to master to achieve their goals at each level. While AI is software, the proposed model keeps a sharp focus on AI adoption by focusing on AI-relevant practices. This will reduce duplication with other software development, organizational development, and digital transformation improvement efforts.

As the gap widens for organizations struggling to turn AI potential into meaningful value, the AI Adoption Maturity Model provides a decisive path forward. It helps organizations build an approach aligned with their business priorities and enables repeatable, scalable success in AI adoption across the enterprise. By emphasizing predictable, evidence-based measurement, the model empowers organizations to realize sustained value from their AI investments, with the flexibility to target the capabilities and practices that deliver the greatest impact in specific domains and provide a clear, actionable roadmap to improve key outcomes. Over time, widespread use of this model will elevate industry-wide capability, driving more consistent, responsible, and effective AI implementation and adoption across sectors.

1 Introduction

Rapid advances in artificial intelligence (AI)—including machine learning (ML), generative AI, and the rise of agentic AI—are forcing organizations to rethink how they operate. Many organizations, and even entire disciplines, find themselves at a turning point where AI and generative AI applications are redefining workflows and challenging existing assumptions of what it means to execute tasks. The discipline of software engineering is a case in point, where tools empowered by generative AI are enabling faster code generation, code summarization, and, in some cases, implementations, despite potential caveats [Ozkaya 2023]. As new technologies become available, software-reliant organizations are exploring where AI can increase efficiency, productivity, and value while reducing costs. This includes embedding AI-enabled capabilities into their internal processes and products to drive innovation. Organizations developing frontier models, such as OpenAI, Google, Microsoft, and Anthropic, continually increase competition by delivering newer versions of their generative AI models with fast-evolving functionalities. The SEI-led National Agenda for Software Engineering had emphasized a future where the existing notions of software development was replaced by one where humans and software are trustworthy collaborators that rapidly evolve systems [Carleton 2021]. Across all disciplines and tasks, the desire for trustworthy human–AI collaboration is quickly accelerating.

Developing AI capabilities, using AI to improve existing practices, or replacing existing capabilities with AI-enabled ones requires many robust practices in addition to those needed to develop the AI models themselves. From having the right data and a strong validation approach to effective human-in-the-loop design and robust infrastructure, all of these enablers add cost and complexity. As a result, many industry and government organizations are struggling to achieve their ambitious value and innovation goals and are seeking guidance on how to adjust their investments and AI adoption strategies to achieve more predictable outcomes that meet their business objectives. Both industry and government organizations need a rigorous approach. The Department of War's (DoW) recently announced Applied AI as a critical technology area, one of many signals that programs need a way to adopt AI quickly, responsibly, and effectively. This critical technology area calls out the need to transform government into an "AI-First" organization to revolutionize decision making and operational efficiency [CTA 2025], which is a shared goal with industry.

1.1 Hype Versus Reality

The harsh reality of early AI experimentation has blunted expected productivity gains and new revenue streams. A 2025 MIT report suggests that despite investments of $30 billion to $40 billion into generative AI, 95 percent of organizations report realizing no returns [Challapally 2025]. AI adoption failures are not limited to the use of generative AI, although the fast-evolving changes around generative AI technologies and agentic AI implementations have certainly contributed to the challenges organizations face.

Even with the confusion in the chaotic race to realize value, some organizations do report gains. A 2023 McKinsey report noted that despite challenges, some organizations have achieved revenue increases and cost reductions in limited areas related to human resources and manufacturing [Chui

2023]. In a 2024 Deloitte report, the number of organizations that stated they were achieving their expected benefits to a "large" or "very large" extent was 18-36 percent [Mittal 2024]. The most recent of these studies, a survey of 2,000 executives conducted by Accenture, shares that about 8 percent of those organizations categorized as front-runners and an additional 7 percent categorized as fast-followers are "AI-reinvention-ready" [Guan 2025]. The successes these organizations report are attributed to strong competency in practices in five core areas, including data and AI strategy; AI platform management; responsible AI; talent and data management; and governance [Guan 2025]. And lastly, the findings of the 2025 Artificial Intelligence Index Report by Stanford University confirm that there are increased investments in, optimism for, and accessibility of AI [Malej 2025].

There is no doubt that a gap exists between the expected magnitude of value from AI technologies and current reality. When organizations fail to see immediate return on investment (ROI) or other forms of value from a technology investment, the cause often is not the technology itself—but a mix of mismatched expectations, misaligned applications, and poorly executed or untested adoption practices. Failures arise when organizations expect the technology to be a "silver bullet" that provides payoffs in a very short amount of time. And we fail to learn from history: technology adoption succeeds when intentional investments are made in building core capabilities that are designed for the context of use and aligned with business goals. Empirical data is consistent: organizations are continuing to invest in AI and do see future business gains. However, realizing the anticipated value of that investment remains unpredictable.

1.2 Why Organizations Fail at AI Adoption

Many organizations have rushed to adopt AI technologies to keep pace with competitors, without objectively evaluating their existing processes and business goals or determining how to align desired outcomes with appropriate AI solutions. While organizations' unique characteristics contribute to many different reasons for failure, studies point to several common root causes. These include not having the right data to support AI initiatives (along with a lack of overall investment in core data management capabilities) [Lewis 2021], a failure to understand user needs, not considering whether a problem would best be solved with AI, a shortage of talent, and, most critically, underestimating the time commitment [Ryseff 2024].

Businesses that are applying AI technologies in the hope that they will provide improved business outcomes need guidance, particularly those businesses that rely heavily on software. Conclusive judgements of success or failure require organizations to identify feasible use cases and define an appropriate scope for AI implementations, use this input to consider the concrete improvements they expect to observe, and then map that progress against an ROI measure. Data collected from the SEI's ongoing empirical research reveal that many organizations struggle to measure and scope a baseline for assessing their AI efforts, both in limited use cases and in the larger context of AI adoption. The inability to effectively scope and baseline leads to gaps in realizing value and makes it difficult to prioritize AI initiatives and identify and mature the key practice areas.

AI adoption requires strong capabilities in organizational change management and system lifecycle engineering. Organizations already struggling with key technology elements essential to AI—such as data management, software competencies, security, and infrastructure—struggle the most as they build AI adoption onto an already shaky foundation. A recurring root cause of failure is technical debt in the AI infrastructure [Kruchten 2019], which manifests as an accumulation of gaps, trade-offs, short-cuts, and lags in compute, networking, data management, security, and talent that compound as companies rush to deploy AI [Cisco 2025]. Avoiding making key technical decisions quickly snowballs into rising costs and barriers in successful AI deployments.

To achieve ROI, organizations must adopt practices that help them determine which workflows can benefit from AI, identify the supporting tools and technologies required, address skill gaps, prioritize critical practices to support AI initiatives, and ensure technical work aligns with business objectives. Organizations need to prioritize predictable and measurable outcomes with intentional goals over bringing in technology in an effort to "keep up." Organizations need a reliable instrument to guide them through the process.

1.3 Driving Maturity Through Predictable, Repeatable Measurement and Analysis

The SEI conducted over two dozen interviews with executives leading AI initiatives, and the top two common challenges voiced in these discussions were 1) a lack of clarity in what to baseline and measure against for improvements when adopting AI and 2) the inability to successfully scale solutions across the enterprise. Both challenges call for recentering the AI adoption journey around institutionalizing key practices with concrete goals and organizational change management practices, implementing intentional communication strategies, and realigning organizational culture as needed.

Institutionalization requires organizations to objectively express what needs to be accomplished and determine the key activities that need to be methodologically completed to do so. They must also focus on what needs to be measured and with what technique, as the goal of institutionalizing practice is to create agile, nimble processes which can be easily implemented, used by the workforce, and assessed effectively. However, organizations need to first focus on baselining their existing business processes and identifying areas of challenge and opportunity. They can then map where AI may remove some of the identified barriers. For example, many off-the-shelf AI-enabled human resource solutions streamline employee management by reducing tool switching and centralizing common tasks such as time-off requests, pay slip retrieval, benefits management, and onboarding. This example shows that organizations may choose either an off-the-shelf solution or a custom-designed one. Each option involves different activities and ROI considerations, which well-institutionalized AI capabilities can effectively support.

When adopting AI, organizations are sometimes narrowly focused on the successful implementation of their use case, without enough consideration given to the success of the overall investment and an understanding of the predictability and scalability of their processes and practices. This problem is at the heart of many of the challenges in baselining and scaling. Even when the predictions about an AI implementation successfully match its outcomes, this accuracy of prediction alone does not reflect the usefulness of the AI implementation to the system, an increased ROI, or the ability to create a scaled practice.

An important first step in addressing AI adoption problems is to treat adoption as the orchestration of process steps which can be controlled, measured, and improved against key intermediary goals along the adoption journey. Similar to any process initiative, from manufacturing to software development, AI adoption and the execution of AI initiatives involve a sequence of tasks that, when properly performed, produces the desired outcome and can be repetitively applied with appropriate use cases and clear outcome measures.

1.4 AI Is Software

Developing and sustaining AI-enabled systems and workflows share much in common with building and maintaining traditional software—yet AI also introduces challenges that organizations cannot ignore [Ozkaya 2021]. AI components often carry a higher margin of error due to the inherent uncertainty of predictive algorithms and the probabilistic nature of generative models. As a result, data management, human-in-the-loop validation, responsible AI practices, and robust security must take on heightened importance, especially as AI systems introduce new attack surfaces.

Similarly, engineering agentic AI systems and workflows that exhibit goal-directed autonomy, reasoning, and continuous evolution add new considerations at the intersection of design, development, and operation. Agentic AI systems need to orchestrate multiple interacting agents, humans, and other tools, requiring rigorous system-level engineering to ensure robustness, safety, observability, and other critical qualities. Still, it is essential to remember that foundational software engineering and cybersecurity practices remain directly relevant. Strong engineering practices—whether developed in-house or accessed through trusted partners—are key enablers of successful AI adoption. AI-enabled systems, whether they are composed of traditional ML components, rely on generative AI services, or orchestrate complex autonomous agentic interactions, are software systems.

Successful AI-enabled systems must be iteratively designed, built, tested, and continuously maintained with engineering discipline. Users need to be confident that the engineering capabilities are sufficient to integrate, test, and monitor AI components as well as manage the needed data. Additionally, existing technologies and infrastructure in the technology stack must be updated in a way that ensures continued operations. Application of certain traditional software and system engineering practices take center stage in developing AI-enabled systems. The following are examples:

- Engineering teams need to architect AI systems for inherent uncertainty in their components, data, models, and output, especially when incorporating generative AI.
- The user experience with AI systems is dynamic. Interfaces must clearly show what the system is doing (i.e., turn-taking), how it generates outputs (i.e., data sources), and when it’s not behaving as expected.
- Engineering teams need to account for different rhythms of change, including change in data, models, systems, and business.

- Verifying, validating, and securing AI systems need to account for ambiguity as well as increased attack surface due to frequently changing data and the underlying nature of models.

The subtle challenge is that AI-enabled systems behave like conventional systems—right up until they don't. This is the point where traditional assumptions break down and where organizations need clarity on how to adapt. The AI Adoption Maturity Model provides this clarity. It highlights the specific practices that must be strengthened or added to existing workflows, helping organizations bridge the gap between conventional software engineering and the unique demands of AI-enabled systems.

The rest of this document provides additional context and describes the Adoption Maturity Model itself: Section 2 summarizes the current state of the practice in AI maturity models, Section 3 describes the proposed model's five levels of AI adoption maturity and their specific characteristics, and Section 4 reviews the next steps to be taken before the model is fully available in Spring 2026.

2 State of the Practice in AI Maturity

AI-adoption problems can be addressed using techniques similar to those used by the software engineering community to improve maturity in the software development process [Humprey 1988]. Identifying key practices and supporting processes needed to apply them with discipline is fundamental to the successful integration of AI into business methods and products and developing AI models and capabilities. The ability to define, implement, perform, monitor, and evolve this set of practices reliably, repeatedly, and consistently will result in capabilities that enable predictable AI adoption. This key idea led the SEI and Accenture to develop the AI Adoption Maturity Model.

2.1 Review of AI Maturity Modeling Efforts

The SEI reviewed publications to develop a comprehensive overview of current AI maturity models and their research and practices, with particular attention to those addressing or referencing generative AI. A search of peer-reviewed journals was conducted using the keywords *AI maturity framework, AI maturity assessment, AI maturity model, AI readiness assessment, and AI capability model.* The search identified 57 sources that were promising enough for a detailed review. Additional expert judgment and internet searches resulted in 58 more sources from grey literature, including models proposed by government and commercial organizations, including consulting companies. A source on agentic AI was included as part of this effort as agentic AI use increased and more information about agentic AI maturity emerged. Any items that were determined to be intended for marketing, with little to no technical detail, were excluded.

The publications deemed relevant to this work described a range of AI-related maturity models and model development activities. The level of detail and rigor varies. Some sources proposed full AI maturity models but did not show evidence of use in practice, while others were research efforts that were limited in scope. Other publications stopped short of defining a model, but included useful elements, such as lists of questions to help organizations begin thinking about the practices needed to drive successful AI initiatives.

Out of the total 115 sources initially selected, 68 explicitly contained some elements of a maturity model, while the rest were high-level discussions about AI maturity and adoption without an explicit model. The focus of the models the publications described varied: 40 focused on AI in general, 7 on generative AI, 5 on responsible AI, and the rest on a variety of very specific topics, such as blockchain and industrial AI. Fourteen models were eliminated due to a lack of detail, leaving 54 models for analysis. Eighteen of these models were excluded as they had only limited discussions. The remaining 36 were studied in detail. Only 26 of these included practice areas, but they did not include components to assist institutionalization. The maturity models the SEI studied mostly focused on common capability areas related to ethics, responsible AI, strategy, innovation, talent, skillsets, people, governance, organization, technology, and data.

The review findings suggest that while there are a number of efforts to develop maturity models that provide good insights into the areas and capabilities that organizations need in order to successfully adopt AI, these models share common drawbacks, including

- *Lack of a clear measurement approach to assess maturity.* Existing approaches to AI maturity provide high-level guidance on areas that organizations need to focus on yet do not enumerate concrete activities and actions to drive predictable institutionalization. This leaves a lot of room for varying interpretations.
- *Inability to incorporate and address emerging areas.* As AI technology quickly evolves, organizations find themselves having to establish new practice areas. For example, AI architecting is taking center stage as agentic AI implementation increases and balancing human-in-the-loop approaches with autonomy is becoming a key part of AI solutions. Existing maturity model efforts do not address the impact that fast evolving technologies have on practice areas.
- *Lack of clarity about the sources of the practice areas identified.* While some areas of capability are commonly agreed upon, such as strong data management practices, there are also a wide variety of inconsistent practices included in the maturity models. In some cases, this results in a lack of clarity on the recommended areas of focus, and in other cases it provides high-level guidance without action.

All of the AI maturity guidance available faces the same challenge: limited evidence of real-world value and difficulty staying relevant as technology rapidly evolves. In this rapidly evolving technology climate, organizations also need to be cognizant of the increasing number of standards and guidance to ensure safety, security, and privacy when adopting AI and leading their organizational AI transformation charters. However, organizations that become overwhelmed with the amount of information available and take no action are at a significant and potentially critical disadvantage.

2.2 Building on a History of Maturity Modeling

The Capability Maturity Model (CMM) framework was developed by the SEI when the software industry was struggling with delivering quality software [Humprey 1988]. From 1990-2013, the SEI developed a series of capability maturity models, beginning with the CMM for Software in 1991 [Paulk 1993, Paulk 1997], the People CMM in 1995 [Curtis 1996, Curtis 2009], and then the CMMI models from 2001-2013 [Chrissis 2011]. The SEI extended its expertise with related maturity model efforts including the influential CERT Resilience Maturity Model (CERT-RMM) [Caralli 2016]. More recently, it co-developed the Cybersecurity Maturity Model Certification (CMMC) [Gardner 2021].

Technology and society have changed considerably since the earlier of these models were developed. An increase in attacks on software systems and other threats have led to an increased focus on security, resilience, and risk management. In addition, generative AI and other software technologies have dramatically changed the landscape of organizational risk and opportunity. The AI Adoption Maturity Model, which is being developed by the SEI with partnership from Accenture, builds on foundational maturity and capability concepts and is also expanded to help organizations make more effective use of AI technologies. Our new modeling efforts are intended to help organizations thrive in this new threat and technology landscape and are focused on getting ahead of technology changes.

3 The AI Adoption Maturity Model

The AI Adoption Maturity Model is based on the fundamental idea that organizations must be able to do the following to achieve winning AI implementations:

- align their desired objectives for AI with their organizational goals and vision
- understand the current state of their capabilities as relevant to AI
- establish a list of required improvement and implementation actions, in order of priority, against their desired state
- identify dependencies across actions and their priorities
- commit resources and time to enable the successful execution of the roadmap

A maturity model provides a framework that helps assess an organization's or unit's ability to perform and sustain specific technical practices in order to achieve its goals. Maturity models outline stages of development and organizational competence, with each stage representing a higher level of organizational capability in a specific area. As such they highlight key critical practice areas and provide a roadmap for improvement. A maturity model is as effective as the robust data and theory it relies on for the development of its structure and for the evidence of its use in practice.

To sustain and develop capabilities relevant to AI, a maturity model needs to focus on institutionalization. This is achieved through concrete descriptions of capabilities, their owners, the processes and practices around their execution, and established approaches for measurement and analysis.

However, the critical reality about AI is that technology not only moves fast, but it sometimes moves in unpredictable directions. Disruptive technologies continue to affect the evolution of both AI and the technologies that support AI. To provide value while taking into consideration these challenges, the AI Adoption Maturity Model does the following:

- focuses on repeatable, predictable measurement and analysis to drive maturity in AI practice
- provides flexibility to create customized views tailored to different domains, organization types, sizes, and contexts
- targets development of a roadmap as a key outcome in addition to an assessment of the maturity of capabilities
- includes AI-relevant capability areas designed to help an organization adapt to changes in the technology landscape
- provides logical groupings of these key AI-relevant capability areas through dimensions

3.1 Building Blocks of the AI Adoption Maturity Model

The basic structural elements of the AI Adoption model and their relationships are shown in Figure 1

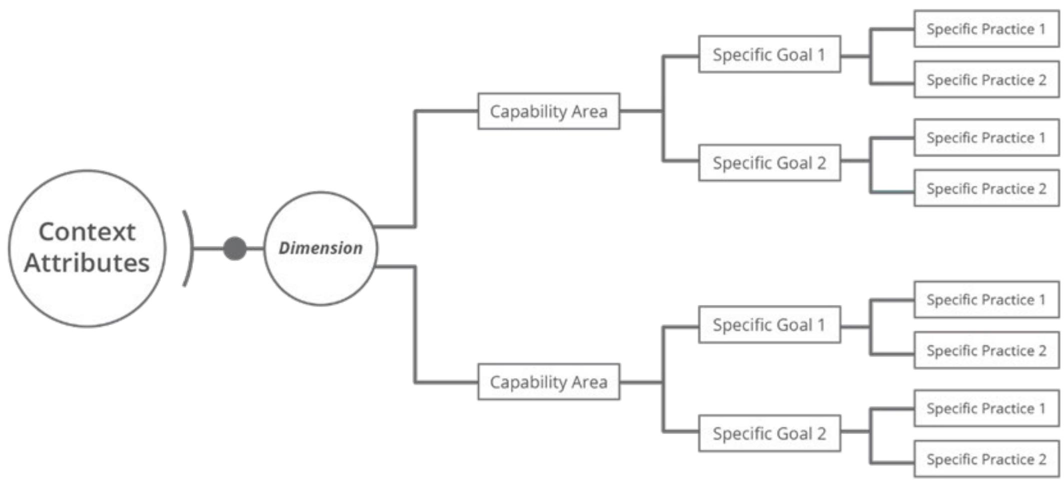

Figure 1: Building Blocks of the AI Adoption Maturity Model

These building blocks are defined below.

- *Context Attributes* – elements that characterize the unit or organization being reviewed based on its organizational, technology, and AI functions to most appropriately assess its current and target capabilities.
- *Dimensions* – the core foundational aspects of AI adoption and maturity which group capability areas. Each dimension along with its capability areas need to address *maturity indicators*—characteristics to demonstrate maturity in addition to implementing practices.
- *Capability Areas* – logical groupings of related goals that together improve an organization's performance in specific skills and activities.
- *Specific Goals* – high-level statements of the outcome to be achieved by the effective implementation of a group of practices within a capability area.
- *Specific Practices* – activities that are considered important to achieve a goal.

3.2 Eight Core Dimensions of Capability

Dimensions are the core foundational aspect of AI adoption and maturity and are used to group the capability areas of the AI Adoption Maturity Model. Each dimension includes maturity indicators or characteristics that demonstrate maturity in relation to implementing practices. The AI Adoption Maturity Model includes eight key dimensions, which are discussed in the following sections.

3.2.1 Organizational Change Dimensions

The four *Organizational Change Dimensions,* described in detail below, focus on business, workforce, and risk-related capability areas.

Organizational Strategy focuses on how effectively an organization positions itself to deliver business value through AI, whether its adoption is limited in scope or integrated broadly across the enterprise. It includes the capability areas which focus on *AI Strategy*, *Aligned Organization Structure*, *Organizational Values and Principles, AI Partnerships and Procurement,* and *Organizational Capability Management.*

Workforce & Culture focuses on capability areas to grow an AI-literate organization through its people, including the training and development that they will need to succeed in integrating, developing, and/or using AI solutions appropriately. It includes the capability areas *Skills and Training*, *Organizational Culture*, and *AI-Ready Workforce Development*.

Workflow Re-engineering emphasizes the changes in workflows and the experimentation that will be needed to use AI to transform key business practices, assist and augment people's work, and take proper advantage of AI in the business, including following sound measurement and analysis practices. This dimension is core to innovation with AI and includes the capability areas *Experimentation*, *Technology Evolution, Business Workflow Innovation*, *Human-in-the-Loop Automation,* and *Measurement and Analysis*.

Risk & Governance focuses on capability areas that cover key standards, policies, and practices to establish and follow to manage AI and related risks both from the perspective of the organization and end users. It includes the capability areas such as *Compliance, Policy, and Standards*, *Risk Management*, and *Responsible AI*.

3.2.2 Lifecycle Engineering Dimensions

Lifecycle Engineering Dimensions focus on the technical and engineering aspects of AI adoption related to AI-enabled systems and workflows. As is the case with the Organizational Change Dimensions, system Lifecycle Engineering Dimensions build on existing software engineering and cybersecurity practices.

Data focuses on maintaining relevant, high-quality, and secure data practices that directly strengthen the value and alignment of AI initiatives. This dimension includes the capability areas such as *Data Management, Data Acquisition and Quality*, and *Data Access, Security, and Privacy*.

Engineering focuses on the key capabilities required to transform AI use cases into production-ready solutions including integrating, developing, and testing AI components and solutions. Example capability areas include *AI Architecting*, *Technology Integration, Transparency and Explainability*, and *Test and Evaluation*.

Operations & Sustainment brings together capability areas essential for collecting production data, monitoring systems for success, and establishing feedback loops for effective long-term sustainment and evolution. Capability areas underscore *Monitoring* to enable predictable, value-added AI adoption, *Production Data Collection*, and *Model Management,* in addition to the ability to build relevant pipelines to support AI initiatives referred to as **Ops*.

Ecosystem focuses on the capabilities needed to properly use and keep up to date on external AI technologies and required infrastructure to ensure the efficiency and scalability of AI solutions.

Capabilities areas include *Supply Chain Management*, *Technology Infrastructure,* and *Deployment at Scale*.

To sustain an organization's focus on AI adoption, the content of each of these capability areas is focused on identifying AI-relevant practices only, though there may be merit and efficiencies to be gained from a broader focus on more general aspects of process maturity across software and cybersecurity.

Figure 2 shows all of the capability areas associated with the dimensions.

Figure 2: Dimensions and Capability Areas

3.3 An Adaptable Maturity Assessment Framework

An AI adoption model needs to cover a minimal set of fundamental capability areas while providing flexibility to adapt to changes in technology evolution, recommended practices, and different organizational contexts. The AI Adoption Maturity Model accomplishes this scoping and adaptability through a set of *context attributes*. Context attributes provide flexibility in the use of the model, guiding the scope of capability areas, specific goals, and practices to assess and steer the organization to those capabilities that matter most to align AI initiatives with their goals. In addition, these context attributes will provide future flexibility to the application of the model and enable different variations in improvement approaches pursued from an empirical basis.

Understanding an organization's AI development approach is one way of scoping with context attributes. This knowledge forecasts how fully the organization will directly pursue its AI initiatives and thus helps in determining what capabilities are most essential to its situation. An organization that aims to fine tune and customize AI models will need different competencies than one that will buy or acquire them instead. An organization that aims to build its agentic AI systems in-house will need different competencies than one that will have other organizations build these systems for them. A number of context attributes focused on organizational functions, technology functions, and AI functions help scope the assessment and subsequent roadmap activity. Table 1 summarizes some of these attributes.

Table 1: Representative Context Attributes

	Attributes	**Description**	**Example Variables**
Organizational Functions	Organization Footprint	The organization/unit's operational footprint, which includes its physical locations, legal entities, and so on	National, Regional, Global
	Domain	Area of operation, such as its industry, products, or customer base	Finance, health care, avionics
	Regulatory posture	Extent and scope of rules, controls, and oversight that an organization is subject to, usually from governmental regulations	Regulated, semi-regulated, un-regulated
	Size	Number of staff relevant to organization/unit	<100, 100-1K, 1K-10K, 10K-100K, >100K
	Planning horizon	Timeline for specific AI adoption goals for the organization/unit	>1 year, 1-3 years, 5years, >5 years
	Scope	Boundary of the AI adoption maturity discussion	Team, business unit, business sector, program, organization
Technology Functions	**Attributes**	**Description**	**Example Variables**
	Software engineering strategy	The extent to which software development supports the organization's business goals	Software organization, only IT functions, software outsourced
	Core technology competency	Structure of the scope and strategy of the organization/unit competencies	Product first, AI first, single core-competency, multi-industry
AI Functions	**Attributes**	**Description**	**Example Variables**
	AI development approach	The extent to which the organization/unit will develop AI initiatives	Build, buy, customize, hybrid
	AI risk profile	The degree of risk the organization/unit is comfortable assuming	High, medium, low
	AI business goal	The archetype that characterizes the organization's approach to AI adoption	Thought leader/AI innovator, early adopter/innovator with AI, follower, laggard
	AI staffing	The organization/unit approach to staffing AI initiatives	In-house, contracted, none

The maturity model is designed to evolve as the technology landscape and organization scope change. The notion of context attributes is one mechanism used to accomplish this by focusing on capability areas that are most relevant to the characteristics of the organization. In addition, the context attributes give the maturity assessors a mechanism to include other relevant practices based on the characteristics of the organizational and technology landscape.

3.4 Progressing in Maturity: AI Adoption Maturity Model Levels

The AI Adoption Maturity Model characterizes the AI adoption journey into one of five maturity steps to support an organization's goals. Figure 3 shows the five levels of AI adoption maturity.

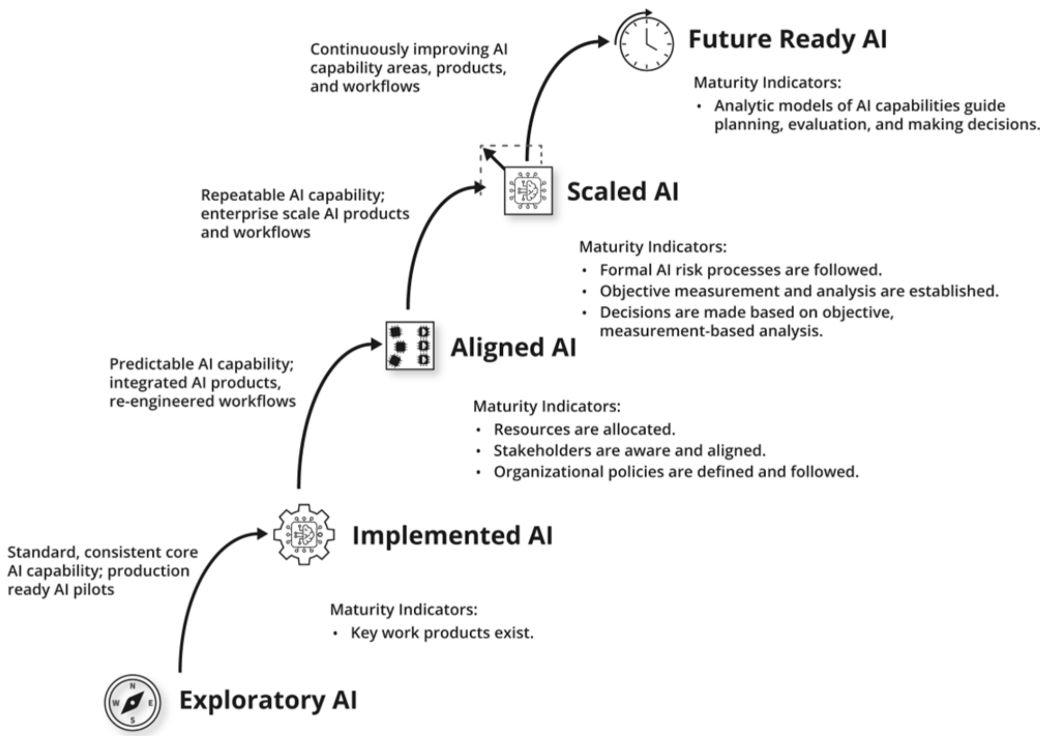

Figure 3: The Five Levels of AI Adoption Maturity

The AI Adoption Maturity Model can help an organization choose what level of maturity to pursue and how they can manage their roadmap. Levels are an abstraction useful for executives and align with an overall AI strategy, while continuous representation by dimensions and capability areas drive actionable prioritization for a roadmap.

Scope is highly important. The scope of adoption maturity does not always apply at the organization level out of the gate. As the scope of its AI adoption, an organization may choose to focus on a single implementation of AI, a business unit, a sector, or a product line. The levels in this maturity model apply to any of the organizational scopes as long as the business goals and scope are clear and are reflected in the roadmap accordingly. We use the term "unit" to refer to the people who are managing the scope of AI adoption and being assessed for maturity. We reserve the term organization for the larger context—creating a roadmap for the broader organizational adoption is a larger scope and has different drivers than starting AI adoption with a more limited scope.

The five levels of maturity are as follows:

1. **Exploratory AI:** At this level, the unit is learning about AI capabilities with regard to the identified scope, risks, and adoption. The goal at this level is to ground AI adoption in the

unit's context, culture, and objectives. Exploration has limited scope and is the start of transformation of a limited set of use cases to establish minimal standard practices. AI-enabled systems and workflows may exist in both sandbox and non-reported production environments.

2. **Implemented AI:** The unit is implementing AI-enabled systems and workflows resulting in some AI-enabled systems and workflows in production. At this level priorities include making sure appropriate practices are in place as the production of AI-enabled systems and workflows carries inherent risks. The unit is on the path for successful AI adoption demonstrating the realization of some use cases.
3. **Aligned AI:** At this level the unit has accomplished the process of AI integration with potential positive impacts on core business services and ROI. Appropriate workflows are identified and managed consistently with AI applications.
4. **Scaled AI:** AI has been integrated into operations, with predictable performance, across a diverse enterprise or community. The entity as well as the organization can replicate AI initiatives to new use cases with relative ease and predictably.
5. **Future Ready AI:** The unit has a track record of success powered by AI. The unit can consistently scale AI initiatives globally and seeks to maintain and evolve their success with AI over a long planning horizon. AI initiatives drive successful and predictable innovations regularly.

These levels of maturity were selected after reviewing the challenges and successes that organizations are realizing; studying the existing software, AI, and cybersecurity maturity landscape; and empirically collecting experiences from the field. A key insight these efforts revealed is that AI adoption maturity, even at low levels of maturity, needs to have an intentional focus. While AI capabilities may be ad hoc or inconsistent and lack practices or measurement, this is not what defines maturity. Intentionally defining scope against goals and investing in key practice areas for that scope is what defines maturity. As such, the AI Adoption Maturity Model is developed to drive success at any level by staging and recommending key capability areas to invest in. Advancing in maturity, therefore, is achieved through realizing the key capabilities across the eight dimensions essential for the respective level. In addition, advancing maturity for each dimension implies realizing the maturity indicators, which include key attributes of established capability area maturity such as resource allocation, relevant governance, and evidence of defining and managing key activities.

While many factors influence maturity transitions, the core objective is to build predictable and repeatable capabilities, enabling organizations to adapt to rapidly evolving AI technologies with measurable practices that reliably produce expected business outcomes.

3.4.1 Exploratory AI

Exploratory AI maturity is defined by experimentation, and the first step of learning and assessing the scope at which AI should be applied. Exploration is intuitively how several organizations reacted when first versions of generative AI models were introduced. Several responded and continue to respond by defining use cases where generative AI could be applied, starting sandbox initiatives at their organizations to see if they could implement simple use cases, and/or holding hackathons to expand the aperture of use cases to select from.

At the *Exploratory level* the unit is assessing use cases for AI applicability and resources for feasibility. The unit may not have an AI strategy beyond a goal to understand where AI may be effective in the organization. A key goal is learning from experimentation.

An entity can achieve the Exploratory AI level by focusing on two key capability areas: *experimentation* and *skills and training*. Key actions therefore to accomplish the Exploratory AI level of maturity include the following:

- Establishing an ability to develop, select from, and assess use cases is key to experimentation. The *experimentation* capability area emphasizes a unit's ability to design, conduct, evaluate, and learn from controlled pilots of AI technologies.
- Developing *skills and training* capability at the initial Exploratory AI level of maturity is an essential component for AI adoption and a reflection of the reality of AI technology. Generative AI, more so than other forms of AI, comes with its own opportunities and challenges. The inherit uncertainty in the structure of the model, the inbuilt variation of correct outcomes, the potential for a larger security attack surface, and the organization's data being potentially compromised necessitates basic skills and training to be available from the start.

Today many routine tools for running business as usual are being revised to include many generative AI-based features in addition to other forms of AI components. Across the board, those organizations that see success reflect having access to talent. On the flip side, access to skillsets is among the top challenges for organizations in their AI adoption experiences. Building skills and training out of the gate are an essential component of Exploratory AI.

An essential component of achieving Exploratory AI level is learning how to scope and resource an AI project in a predictable way. While many organizations know to start with experimentation, they do not include an explicit validation step for comparative and predictable outcomes. Effective exploration sets the foundation of effective maturity improvement.

3.4.2 Implemented AI

Implemented AI maturity is reached when a unit starts deploying, even in a limited scope, core AI-enabled features, internally for organizational use or externally to their customers. Units targeting Implemented AI are deploying AI-enabled systems, workflows, and products in production, either involving production data or solutions that the unit is using to support business processes.

Examples of drivers that enable maturity at the level of Implemented AI are characterized by the following:

- Relevance of the use cases are piloted in the field through *Operational Experimentation*. This will require several key capabilities to be developed, including *Data Acquisition and Quality, Data Management,* and *Data Access, Security, and Privacy*.

- Ability to deliver AI-enabled features, even at a limited scope, will require engineering competencies, bringing *AI Architecting* capability to the forefront.
- *Measurement and Analysis* will need to be mastered to understand the resources needed for the AI initiative as well as assessing the effectiveness of the capabilities.
- AI implementations are known and understood more broadly and are governed and risk managed appropriately.

The following are key actions to implement to ensure successful Implemented AI maturity and establish a path forward to effectively focus on the innovation and workflow re-engineering required to achieve the next level of Aligned AI maturity:

- The unit will need to master all aspects of data management and governance. Established organizations will likely have existing practices around data management. These can be extended for relevance to AI. However, AI adoption, especially of generative AI, creates a smaller scope for new requirements around data management, establishing pipelines for preparing data, governance around data, and security aspects that all apply evenly.
- Similarly, capabilities around AI model management will need to be established. A core challenge with AI adoption is keeping up with the pace of evolution. Without defined processes and practices which drive model evolution, organizations find themselves following ad hoc practices which result in unnecessary resources spent when new versions of AI models are released. Experiences of degrading system outcomes because of degrading model performance of newer versions are common. Units will need to establish controlled assessment practices to be able to make decisions with an empirical basis.
- Organizational values and principles should remain consistent. AI adoption should adapt to organizational values and principles, ensuring that goals are communicated clearly and stakeholders are aware of roles and responsibilities.

A unit which has achieved Implemented AI level across relevant dimensions demonstrates successful AI-enabled system and workflow deployment and is ready to build on this foundation to achieve organizational transformation through AI.

3.4.3 Aligned AI

The Aligned AI level is characterized by workflow re-engineering and a strategic focus on innovation. At this level, the unit has achieved the key practices to develop, monitor, and resource AI initiatives across business units and other functions. The unit has fundamental business practices in place (e.g., workforce, data, measurement and analysis, governance, risk management, and engineering capabilities) to challenge reliably existing approaches within either their internal processes or designed and delivered products to take advantage of AI through aligning AI initiatives along with other tools.

While at the Aligned AI level units are realizing changed workflows and establishing key capabilities to achieve repeatable success, they have not yet scaled AI across the organization. The unit needs to demonstrate a key strategic initiative to position itself for successful Scaled AI maturity. This is where the core infrastructure and business investments occur and build upon established measurement and analysis and monitoring practices for informed decision making.

Key actions to take to achieve maturity at the level of Aligned AI and be on the path for Scaled AI include the following:

- Demonstrate competency in *Workflow Re-engineering* as part of Aligned AI level; this is also the core of the Scaled AI level.
- Invest in an *AI-Ready Workforce* development capability.
- Ensure core processes and practices for managing AI risk are supported by defined processes, including *Responsible AI and Human in-the-loop Automation, Transparency and Explainability and Compliance,* and *Policy and Standards.*
- *AI Strategy* and *Aligned Organizational Structure* capabilities will identify the vision, responsibilities, and goals of AI initiatives which will be the foundation for success at scale.

Maturity indicators for each dimension will need to demonstrate objective evaluation of policy adherence, descriptions of how exceptions are addressed, high-level management involvement and ownership of risks, resource allocations, and commitment to removing barriers to adoption.

3.4.4 Scaled AI

As an entity moves through the Exploratory AI, Implemented AI, and Aligned AI levels implementing key capability areas and pursuing maturity indicators of governed, measured, and predictable practices for each applicable dimension, it will have set itself up on the right path to achieve repeatable and scalable AI initiatives.

A complexity that AI adoption brings is the increasing need to manage the supply chain and AI ecosystems through partnerships. An organization at this level will have already demonstrated how to select partners and assess their AI capabilities against the organization's expectations.

The key indicator of success at the Scaled AI level is the ability to objectively define what AI capabilities are needed to scale AI across the organization. As the organization will have already established mechanisms to identify and allocate resources and responsibilities, the focus of Scaled AI will be replicating these initiatives by employing gained knowledge and focusing on shared resources to optimize costs.

Key maturity indicators to advance to the Future Ready AI level include the following:

- Demonstrate *Deployment at Scale* of capabilities where the business, technical, and infrastructure requirements are aligned to support AI-initiatives across the organization.
- Strategic decisions include focused discussions on how AI can support increased value to both the organization and to its workforce.
- Shared artifacts define the success of AI initiatives, not only those limited to select AI functions which may have been the case when the organization was at the Implemented AI or Aligned AI stages.
- At the Scaled AI level, organizations have also established optimal *AI Ready Workforce* development initiatives and are gaining benefits from these initiatives, including sharing talent, tools, procedures, and practices across initiatives as well.

- The organization will also need to begin mastering *Organizational Capability Management,* demonstrating the use of analytic models, and continuing data collection and analysis to drive decision making.

3.4.5 Future Ready AI

Achieving the Future Ready AI level means that the organization has made the intentional decision to be an AI-first company where AI is embedded in the organization's culture, the organization knows how to take advantage of old and new advances rapidly, and the organization continues to innovate with AI while augmenting it with other relevant tools and approaches.

A Future Ready AI organization can assess new AI innovations for their relevance, usefulness, and timeliness and compare them to the existing AI initiatives with concrete criteria, employing a predictive capability to more accurately forecast their implications for enhanced organizational performance and business objectives. The organization already has foundational capabilities and processes and practices to resource new initiatives and can evolve them as needed.

Organizations (or units) at Future Ready AI level have achieved effective implementation of the *Organizational Capability Management* capability of the Organizational Structure dimension and demonstrate the following:

- Improvements are identified proactively using established analytic techniques (e.g. causal analysis, statistical methods, simulations).
- Existing processes and practices and analytic techniques are utilized to assess new technical advancements in the AI landscape and move initiatives quickly through the Exploratory, Implemented, Aligned, and Scaled levels, effectively and with planned and monitored outcomes.
- Community outreach initiatives are established where they support and influence the AI practitioner, tooling, and research community in innovations.
- People know their responsibilities, know where the resources are, and, when changes occur, are not perplexed and can rely on the established practices and analytic techniques to make informed assessments about the relevance of the AI technology and its expected outcomes.

As an example, organizations today struggle to assess the relevance of agentic AI and where to employ it. Agentic AI is not a new form of AI—though it is often mentioned in conjunction with advancements of generative AI and foundation models. Agentic AI brings together elements in the AI landscape to develop goal-oriented, adaptive, self-corrective capabilities to support the development of AI-enabled systems by delegating the decision making to autonomous AI systems. A Future Ready AI organization would be able to rely on its established practices in *Experimentation* to determine use cases relevant to apply agentic AI, rely on existing *Responsible AI and Human-in-the-loop Automation* practices to assess governance risks, and borrow skills and practices from *AI Architecting* to incorporate agentic AI into its system design decision tradeoffs. At the Future Ready AI level, the organization has the required capabilities and maturity indicators to avoid surprises with data and anticipate future changes.

Figure 4 shows the capability areas to achieve each maturity level.

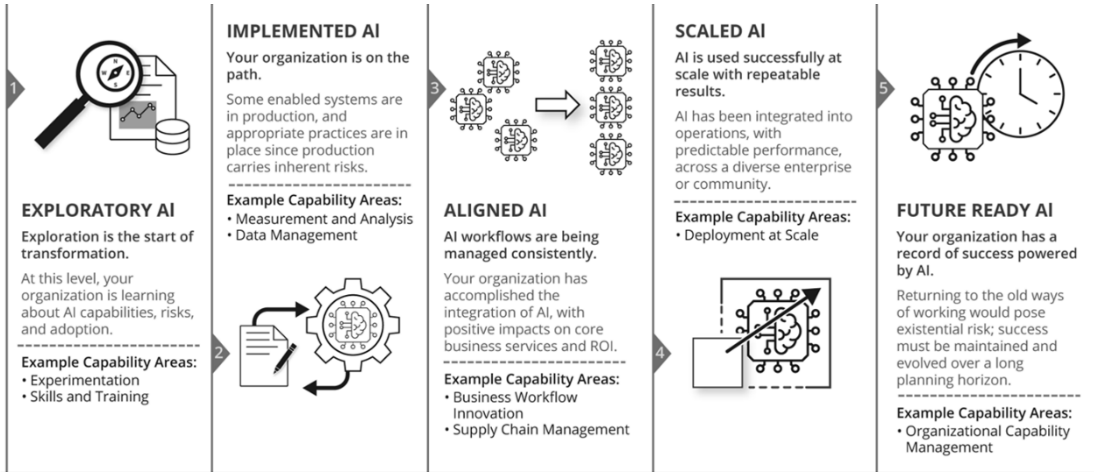

Figure 4: Key Capabilities of AI Adoption Maturity

3.5 Roadmap Development

The goal of this AI Adoption Maturity Model effort is not certification of units or organizations against their use case implementation and competency on each of the dimensions, but to create an actionable roadmap to drive institutional AI adoption.

A roadmap is the natural outcome of assessing the dimensions for maturity by focusing on key capability areas and their dependencies and priorities. The context attributes further provide the necessary scoping to align the investments in core AI adoption capabilities that are needed. A unit would have the ability to align their business goals and accomplish maturity advancement by scoping different initiatives at different levels, if needed. The SEI is currently conducting pilots and will share lessons learned and their mapping to the AI Adoption Maturity Model elements as they are completed.

4 Conclusion and Next Steps

There has been much written about the success and failures of AI and AI adoption. These lessons learned capture individual reflections and perceptions but typically contain limited to no factual data. Such lessons learned provide input toward identifying areas where improvement is needed, but they do not provide actionable steps. The AI Adoption Maturity Model aims to help fill this gap by providing a framework that organizations can use to identify priority areas for immediate attention and improvement. In addition, it provides recommended practices that target key challenge areas, such as baselining and measuring improvements, workflow re-engineering, and managing the AI ecosystem.

Despite the challenges organizations are experiencing in the race to innovate with AI and their mixed initial results, the opportunity exists to realize much of AI's promised potential with the right efforts. History has taught us repeatedly that there is no silver bullet to replace discipline and intention. The AI Adoption Maturity Model is structured to provide guidance exactly where it is needed, with practices that intentionally guide an organization's AI adoption journey with predictable outcomes.

The SEI is currently conducting pilots applying the AI Adoption Maturity Model and would welcome the opportunity to collaborate with organizations who would like to join our early adopter program. Those interested can email info@sei.cmu.edu. The full model will be publicly available in Spring 2026. A much needed outcome of this effort will be benchmarks across organizations providing insights into areas of maturity and actionable best practices.

Bibliography

URLs are valid as of the publication date of this report.

[Caralli 2016]
Caralli, R. A., Allen, J. H., & White, D. W. *CERT Resilience Management Model (CERT-RMM): A maturity model for managing operational resilience.* Addison-Wesley Professional. 2016. ISBN 978-0-13-454506-6

[Carleton 2021]
Carleton, A., Klein, M. H., Robert, J. E., Harper, E., Cunningham, R., de Niz, D., Foreman, J. T., Goodenough, J. B., Herbsleb, J. D., Ozkaya, I., Schmidt D., and Shull, S. *Architecting the Future of Software Engineering: A National Agenda for Software Engineering Research & Development.* Carnegie Mellon University, Software Engineering Institute. 2021. https://www.sei.cmu.edu/documents/1308/2021_014_001_741195.pdf

[Challapally 2025]
Challapally, Aditya et al. *GenAI Divide: State of AI in Business 2025*. MIT NANDA. July 2025. https://mlq.ai/media/quarterly_decks/v0.1_State_of_AI_in_Business_2025_Report.pdf

[Chui 2023]
Chui, Michael et al. *The state of AI in 2023: Generative AI's breakout year*. QuantumBlack, AI by McKinsey. August 1, 2023. https://www.mckinsey.com/capabilities/quantumblack/our-insights/the-state-of-ai-in-2023-generative-AIs-breakout-year

[Chrissis 2011]
Chrissis, M. B., Konrad, M., & Shrum, S. *CMMI: Guidelines for process integration and product improvement* (3rd ed.). Addison-Wesley Professional. 2011.

[Cisco 2025]
Cisco. *Realizing the Value of AI. Cisco AI Readiness Index 2025*. Cisco. 2025. https://www.cisco.com/c/dam/m/en_us/solutions/ai/readiness-index/2025-m10/documents/cisco-ai-readiness-index-2025-realizing-the-value-of-ai.pdf

[Cooper 2024]
Cooper, Robert G. Why AI Projects Fail: Lessons From New Product Development. *IEEE Engineering Management Review*. Volume 52. Number 4. August 2024. Pages 15–21. https://doi.org/10.1109/EMR.2024.3419268

[Curtis 1995]
Curtis, W., Hefley, W. E., & Miller, S. (1995). *People Capability Maturity Model (P-CMM)* (CMU/SEI-95-MM-01). Software Engineering Institute, Carnegie Mellon University. https://www.sei.cmu.edu/documents/1634/1995_008_001_16349.pdf

[Curtis 2009]
Curtis, W., Hefley, W. E., & Miller, S. A. *The People Capability Maturity Model (P-CMM)*, 2nd ed. Addison-Wesley Professional. 2009. ISBN 978-0-321-55390-4

[CTA 2025]
Critical Technology Areas, Office of the Under Secretary of War for Research and Engineering. December 2025. https://www.cto.mil/cta/

[Gardner 2021]
Gardner, D. *Overview of Practices and Processes of the CMMC 1.0 Assessment Guides (CMMC 1.0)* Software Engineering Institute, Carnegie Mellon University. 2021. https://www.sei.cmu.edu/library/overview-of-practices-and-processes-of-the-cmmc-10-assessment-guides-cmmc-10/

[Guan 2025]
Guan, Lan; Ramani, Senthil; & Roussiere, Philippe. 2025. *The front-runners' guide to scaling AI: Lessons from industry leaders*. Accenture. May 6, 2025. https://www.accenture.com/content/dam/accenture/final/accenture-com/document-3/Accenture-Front-Runners-Guide-Scaling-AI-2025-POV.pdf

[Humprey 1988]
Humprey, Watts S. Characterizing the Software Process: A Maturity Framework. *IEEE Software.* March 1988. Volume 5. Issue 2. Pages 73-79. https://doi.org/10.1109/52.2014

[Kruchten 2019]
Kruchten, Philippe; Nord, Robert; & Ozkaya, Ipek. *Managing Technical Debt: Reducing Friction in Software Development*. Addison-Wesley Professional. 2019. ISBN: 978-0135645932.

[Lewis 2021]
Lewis, Grace A.; Bellomo, Stephany; & Ozkaya, Ipek. Characterizing and Detecting Mismatch in Machine-Learning-Enabled Systems. Pages 133-140. In *2021 IEEE/ACM 1st Workshop on AI Engineering - Software Engineering for AI (WAIN)*. 2021. https://doi.org/10.1109/WAIN52551.2021.00028

[Maslej 2025]
Maslej, Nestor et al. *The AI Index 2025 Annual Report*. AI Index Steering Committee, Institute for Human-Centered AI, Stanford University. April 2025. https://hai.stanford.edu/assets/files/hai_ai_index_report_2025.pdf

[Mittal 2024]
Mittal, Nitin et al. 2024. *Now decides next: Getting real about Generative AI*. The Deloitte AI Institute Report. The State of Generative AI in the Enterprise: Quarter Two Report, Deloitte AI Institute. April 2024. https://www.deloitte.com/content/dam/assets-shared/docs/about/2025/quarter-2.pdf

[Ozkaya 2021]
Ozkaya, Ipek. What Is Really Different in Engineering AI-Enabled Systems? *IEEE Software*. Volume 37. Issue 4. July-August 2020. Pages 3–6. https://doi.org/10.1109/MS.2020.2993662

[Ozkaya 2023]
Ozkaya, Ipek. Application of Large Language Models to Software Engineering Tasks: Opportunities, Risks, and Implications. *IEEE Software*. Volume 40. Issue 3. May 1, 2023. Pages 4–8. 10.1109/MS.2023.3248401

[Paulk 1993]
Paulk, M. C., Weber, C. V., Garcia, S. M., Chrissis, M. B., & Bush, M. (1993). *Key Practices of the Capability Maturity Model for Software, Version 1.1* (CMU/SEI-93-TR-025). Software Engineering Institute, Carnegie Mellon University. https://www.sei.cmu.edu/documents/1092/1993_005_001_16211.pdf

[Paulk 1997]
Paulk, M. C., Curtis, B., Chrissis, M. B., & Weber, C. V. (1997). *The Capability Maturity Model: Guidelines for improving the software process* (reprint ed.). Addison-Wesley.

[Ryseff 2024]
Ryseff, James; De Bruhl, Brandon F.; & Newberry, Sydne J. *The Root Causes of Failure for Artificial Intelligence Projects and How They Can Succeed: Avoiding the Anti-Patterns of AI.* RAND Corporation. August 13, 2024. https://www.rand.org/content/dam/rand/pubs/research_reports/RRA2600/RRA2680-1/RAND_RRA2680-1.pdf