

JAMES RYSEFF, BRANDON DE BRUHL, SYDNE J. NEWBERRY

The Root Causes of Failure for Artificial Intelligence Projects and How They Can Succeed

Avoiding the Anti-Patterns of AI

Artificial intelligence (AI) is widely recognized as technology with the potential to have a transformative effect on organizations.¹ Although AI was once reserved for advanced technology companies with the ability to hire top talent and spend millions of dollars, all types of organizations are adopting AI today. Private-sector investment in AI increased 18-fold from 2013 to 2022,² and one survey found that 58 percent of midsize corporations³ had deployed at least one AI model to production.⁴ Similarly, the U.S. Department of Defense (DoD) is spending \$1.8 billion each year on military applications for AI, and DoD leaders have identified AI as one of the most crucial technologies to the future of warfare.⁵

AI is already making impacts across a wide variety of industries. Pharmaceutical companies are using it to accelerate the pace and success rate of drug development.⁶ Retailers, such as Walmart, are deploying AI for predictive analytics so that they know when to restock inventory and how to optimize their end-to-end supply chains.⁷ Finally, in the defense realm, AI is piloting fighter jets,⁸ detecting enemy submarines,⁹ and improving commanders' awareness of the battlefield.¹⁰ These examples demonstrate the relevance of AI to organizations in a variety of industries and for a variety of use cases.

However, despite the promise and hype around AI, many organizations are struggling to deliver working AI applications. One survey found that only 14 percent of organizations responded that they were fully ready to adopt AI, even though 84 percent of business leaders reported that they believe that AI will have a significant impact on their business.¹¹ Managers and directors find themselves under enormous pressure to do something—anything—with AI to demonstrate to their superiors that they are keeping up with the rapid advance of technology.¹² But too many managers have little understanding of how to translate this desire into action. By some estimates, more than 80 percent of AI projects fail.¹³ This is twice the already-high rate of failure in corporate information technology (IT) projects that do not involve AI.¹⁴

SUMMARY

Background

Although leaders widely recognize the importance of artificial intelligence (AI), successfully implementing AI projects remains a serious challenge.^a According to one survey, 84 percent of business leaders responded that they believe that AI will have a significant impact on their business, and 97 percent of business leaders reported that the urgency to deploy AI-powered technologies has increased.^b Despite this, the same survey found that only 14 percent of organizations responded that they were fully ready to integrate AI into their businesses.

By some estimates, more than 80 percent of AI projects fail—twice the rate of failure for information technology projects that do not involve AI.^c Thus, understanding how to translate AI's enormous potential into concrete results remains an urgent challenge. In this report, we document lessons learned from those who have already applied AI/ML so that U.S. Department of Defense leadership and others can avoid these failures or mitigate risks in their planning.

Approach

To investigate why AI projects fail, we interviewed 65 experienced data scientists and engineers. Participants had at least five years of experience building AI/ML models in industry or academia. We selected participants across a variety of company sizes and industries to ensure that these findings would be broadly representative. The output of these interviews is summarized in this analysis.

Takeaways

Our interviews highlighted five leading root causes of the failure of AI projects. First, industry stakeholders often misunderstand—or miscommunicate—what problem needs to be solved using AI. Too often, trained AI models are deployed that have been optimized for the wrong metrics or do not fit into the overall business workflow and context. Second, many AI projects fail because the organization lacks the necessary data to adequately train an effective AI model. Third, in some cases, AI projects fail because the organization focuses more on using the latest and greatest technology than on solving real problems for its intended users. Fourth, organizations might not have adequate infrastructure to manage their data and deploy completed AI models, which increases the likelihood of project failure. Finally, in some cases, AI projects fail because the technology is applied to problems that are too difficult for AI to solve. AI is not a magic wand that can make any challenging problem disappear; in some cases, even the most advanced AI models cannot automate away a difficult task.

Industry Recommendations

To overcome these issues, leaders should consider these five principles for success in AI projects:

- *Ensure that technical staff understand the project purpose and domain context:* Misunderstandings and miscommunications about the intent and purpose of the project are the most common reasons for AI project failure. Ensuring effective interactions between the technologists and the business experts can be the difference between success and failure for an AI project.
- *Choose enduring problems:* AI projects require time and patience to complete. Before they begin any AI project, leaders should be prepared to commit each product team to solving a specific problem for at least a year. If an AI project is not worth such a long-term commitment, it most likely is not worth committing to at all.
- *Focus on the problem, not the technology:* Successful projects are laser-focused on the problem to be solved, not the technology used to solve it. Chasing the latest and greatest advances in AI for their own sake is one of the most frequent pathways to failure.

- *Invest in infrastructure:* Up-front investments in infrastructure to support data governance and model deployment can substantially reduce the time required to complete AI projects and can increase the volume of high-quality data available to train effective AI models.
- *Understand AI's limitations:* Despite all the hype around AI as a technology, AI still has technical limitations that cannot always be overcome. When considering a potential AI project, leaders need to include technical experts to assess the project's feasibility.

Academia Recommendations

To overcome the issues described by our academic interviewees, leaders should consider these two recommendations:

- **Overcome data-collection barriers through partnerships with government:** Partnerships between academia and government agencies could give researchers access to data of the provenance needed for academic research. The federal government should expand its investment in such programs as Data.gov (the U.S. government's open data site) and seek to increase the number of datasets available for research.
- **Expand doctoral programs in data science for practitioners:** Newer academics often feel pressure to focus on research that leads to career success as opposed to research that has the most potential to solve important problems. Computer science and data science program leaders should learn from disciplines, such as international relations, in which practitioner doctoral programs often exist side by side at even the top-ranked universities to provide pathways for the most-advanced researchers to apply their findings to urgent problems.

^a For this project, we focused on the machine learning (ML) branch of AI because that is the technology underpinning most business applications of AI today. This includes AI models trained using supervised learning, unsupervised learning, or reinforcement learning approaches and large language models (LLMs). Projects that simply used pretrained LLMs (sometimes known as *prompt engineering*) were not included in the scope of this work.

^b Cisco AI Readiness Index.

^c Kahn, "Want Your Company's AI Project to Succeed?"

The purpose of this exploratory analysis is to help leaders and managers within all types of organizations who are struggling to understand how to execute AI projects in their organization avoid some of the most common reasons for AI project failures. To do so, we interviewed 65 experienced AI engineers and researchers across a variety of companies and industries, as well as academia. From these interviews, we identified the most frequently reported *anti-patterns* of AI—common responses to recurring problems that are typically ineffective or even counterproductive.¹⁵ We hope to help organizations avoid making these common mistakes and to provide leaders and managers endeavoring to understand AI with practical advice to help them get started.

AI projects have two components: the technology as a platform (i.e., the development, use, and deployment of AI to complete some set of business tasks) and the organization of the project (i.e., the process, struc-

ture, and place in the overall organization). These two elements enable organizations and AI tools to work together to solve pressing business problems.¹⁶

IT-type projects can fail for many reasons not related to the technology itself. For example, projects can fail because of process failures (i.e., flaws in the way the project is executed), interaction failures (i.e., problems with how humans interact with the technology), or expectation failures (i.e., a misalignment in the anticipated value of the project).¹⁷ Breakdowns in any component could result in a project failure, which results in increased costs for the sponsoring enterprise. There is a large body of literature on how IT projects fail. However, AI seems to have different project characteristics, such as costly labor and capital requirements and high algorithm complexity, that make them unlike a traditional information system.¹⁸ The high-profile nature of AI may increase the desire for stakeholders to better understand what drives the risk of IT projects related to AI.

Most prior work on this topic has taken one of two forms. In some cases, an individual data scientist or manager discusses their personal experiences and beliefs about what causes AI projects to fail.¹⁹ In other cases, consulting firms conduct a widespread survey of IT leaders to discuss their experiences with AI.²⁰ For example, McKinsey has conducted an annual survey about AI for several years.²¹ Additionally, one study conducted a systematic literature review and interviews with six experts to explore the factors that might cause general AI projects to fail.²²

Our study differs from this prior work in several ways. First, we focus on the perspective of the individuals building AI applications as opposed to the business leaders of the organization. A bottom-up approach allows us to discuss why AI projects fail from the point of view of the people who intimately understand the specifics of the technology. Second, we conducted semistructured interviews as opposed to relying on multiple-choice or short-answer survey questions. Although the burden of conducting interviews means that the sample size of this study is smaller compared with those of multiple-choice survey studies, this approach allowed us to explore the issues raised in greater nuance and depth. Finally, we conducted substantially more semistructured interviews with experts compared with prior authors who took this approach.

Methods

To gather data for this report, we conducted semistructured interviews with experienced AI practitioners in both industry and academia. During these interviews, we defined the *failure* of an AI project as a project that was perceived to be a failure by the organization. We included both technical failures and business failures within this definition. Each interviewee was asked to discuss the types of failures that they perceived to be the most frequent or impactful and what they believed the root causes of these failures were. We then identified common root causes based on the interview responses. The interviews were conducted between August and December 2023.

The approach taken in this report has strengths and weaknesses. Conducting interviews with open-

ended questions of experienced data scientists and ML engineers allowed us to discover what these professionals believe are the greatest problems and challenges when attempting to execute AI projects. However, because the majority of our interviewees were nonmanagerial engineers instead of business executives, the results may disproportionately reflect the perspective of individuals who do not hold leadership positions. Thus, the results may be skewed toward identifying leadership failures.

Industry Participants

We identified potential industry participants using the LinkedIn Recruiter tool and LinkedIn InMail messages. Potential participants had at least five years of AI/ML experience in industry and job titles that indicated that they were either an individual contributor or a manager in the data science or ML engineering technical disciplines.²³ We selected participants to represent a variety of experiences and backgrounds. In particular, we selected participants from different company sizes (start-ups, large companies, and medium-sized companies) and industries (technology, health care, finance, retail, consulting, and others). Industry participants were offered a \$100 honorarium for agreeing to take part in a 45-minute interview.

A total of 379 potential industry candidates were identified and contacted. Of these, 50 individuals ultimately participated in an interview, representing more than 50 unique organizations.²⁴ Fourteen individuals sent a message declining to participate in the study; these individuals were removed from the candidate pool and had no further contact from the study team.²⁵ Table 1 illustrates the percentages of potential candidates who either participated or declined to participate in the study.

Industry interviews used a consistent battery of questions, which is provided in Appendix A. All interviews were conducted with a promise of anonymity to ensure that participants felt free to speak candidly about their experiences.

Academia Participants

We conducted 15 interviews of academics drawn from convenience samples during conferences and from individuals known to the research team. These interviews ranged across school types (e.g., engineering programs and business schools) and degree levels (e.g., tenure-track researcher, non-tenure-track researcher, graduate student, and undergraduate or research assistant). These interviews used a consistent battery of questions, which is presented in Appendix B. Our interviews were conducted with the promise of anonymity to allow non-tenure-track academic researchers and nonresearcher engineers who support the research efforts to have an opportunity to speak without attribution. Table 2 illustrates the academic candidate response rates.

Findings from Industry Interviews

Across all of the interviews conducted with experienced AI practitioners from industry, five dominant root causes emerged describing why AI projects fail. Overall, interviewees expressed that the most common root cause of failure was the business leadership of the organization misunderstanding how to set the project on a pathway to success. Our interviewees also noted that these types of failures had the most impact on the ultimate outcome of the project compared with the other root causes of failure they discussed.

The other notable root cause of failure identified by interviewees was limitations in the quality and utility of data available to train the AI models. These two root causes were cited spontaneously by more than one-half of the interviewees as the primary reasons that AI projects failed or underperformed.

In addition to the most frequent failure patterns cited, three other root causes were noted by a meaningful number of interviewees.²⁶ First, some interviewees noted the lack of investment in infrastructure to empower the team. Second, some interviewees discussed the difference between the top-down failures caused by leadership and the bottom-up failures caused by individual contributors on the data science team. Finally, some interviewees discussed project

TABLE 1
Industry Candidate Response Rates

Indicators	Candidate Pool	Accepted	Declined
Number of candidates	379	50	14
Percentage	100	13.2	3.7

TABLE 2
Academic Candidate Response Rates

Indicators	Candidate Pool	Accepted	Declined
Number of candidates	37	15	22
Percentage	100	40.5	59.5

failures caused by fundamental limitations in what AI can actually achieve. While these failure patterns were cited less frequently than the two dominant root causes, they each were cited by a one-quarter to one-third of the interview participants.

Leadership-Driven Failures

More than any other type of issue, our interviewees noted that failures driven by the decisions and expectations of the organization's business leadership were far and away the most frequent causes of project failure. Eighty-four percent of our interviewees cited one or more of these root causes as the primary reason that AI projects would fail. These leadership-driven failures took several forms.

Optimizing for the Wrong Business Problem

First, all too often, leadership instructs the data science team to solve the wrong problem with AI. This results in the data science team working hard for months to deliver a trained AI model that makes little impact on the business or organization. In many cases, this is due to a communication breakdown between the data science team and the leaders of the organization.

Few business leaders have a background in data science; consequently, the objectives they set need to be translated by the technical staff into goals that can

be achieved by a trained AI model. In failed projects, either the business leadership does not make themselves available to discuss whether the choices made by the technical team align with their intent, or they do not realize that the metrics measuring the success of the AI model do not truly represent the metrics of success for its intended purpose. For example, business leaders may say that they need an ML algorithm that tells them the price to set for a product—but what they actually need is the price that gives them the greatest profit margin instead of the price that sells the most items. The data science team lacks this business context and therefore might make the wrong assumptions. These kinds of errors often become obvious only after the data science team delivers a completed AI model and attempts to integrate it into day-to-day business operations.

Using Artificial Intelligence to Solve Simple Problems

In other cases, business leaders demand that the technical team apply ML to a problem that does not truly require it. Not every problem is complex enough to require an ML solution: As one interviewee explained, his teams would sometimes be instructed to apply AI techniques to datasets with a handful of dominant characteristics or patterns that could have quickly been captured by a few simple if-then rules. This mismatch can happen for different reasons. In some cases, leaders understand AI only as a buzzword and do not realize that simpler and cheaper solutions are available. In other cases, senior leaders who are far removed from the implementation details demand the use of AI because they are confident that their business area must have complex problems

Many leaders are not prepared for the time and cost of acquiring, cleaning, and exploring their organization's data.

that demand complex solutions. Regardless of the cause, while these types of projects might succeed in a narrow sense, they fail in effect because they were never necessary in the first place.

Overconfidence in Artificial Intelligence

Additionally, many senior leaders have inflated expectations of what AI can be expected to achieve. The rapid advancements and impressive achievements of AI models have generated a wave of hype about the technology. Pitches from salespeople and presentations by AI researchers add to the perception that AI can easily achieve almost anything. In reality, optimizing an AI model for an organization's use case can be more difficult than these presentations make it appear. AI models developed by academic researchers might not work effectively for all of the peculiarities of an organization's business. Many business leaders also do not realize that AI algorithms are inherently probabilistic: Every AI model incorporates some degree of randomness and uncertainty. Business leaders who expect repeatability and certainty can be disappointed when the model fails to live up to their expectations, leading them to lose faith in the AI product and in the data science team.

Underestimating the Time Commitment Needed

Finally, many interviewees (14 of 50) reported finding that senior leaders often underestimated the amount of time that it would take to train an AI model that was effective at solving their business problems. Even when an off-the-shelf AI model is available, it has not been trained on an organization's data and thus it may not be immediately effective in solving the specific business problems. Many leaders are not prepared for the time and cost of acquiring, cleaning, and exploring their organization's data. They expect AI projects to take weeks instead of months to complete, and they wonder why the data science team cannot quickly replicate the fantastic achievements they hear about every day. Even worse, in some organizations, senior leaders rapidly switch their priorities every few weeks or months. In these cases, projects that are in progress can be discarded before they have the opportunity to demonstrate real

results, or completed projects can be ignored because they no longer address what leadership views as the most important priorities of the company. Even when the project is successful, leaders may direct the team to move on prematurely. As one interviewee put it, “Often, models are delivered as 50 percent of what they could have been.”²⁷

Bottom-Up–Driven Failures

In contrast to the top-down failure patterns driven by the organization’s business leadership, many interviewees (16 of 50) noted a different type of failure pattern driven by the data scientists on the team. Technical staff often enjoy pushing the boundaries of the possible and learning new tools and techniques. Consequently, they often look for opportunities to try out newly developed models or frameworks even when older, more-established tools might be a better fit for the business use case. Individual engineers and data scientists also have a strong incentive to build up their experience using the latest technological advancements because these skills are highly desired in the hiring market. AI projects often fail when they focus on the technology being employed instead of focusing on solving real problems for their intended end users. While it is important for an organization to experiment with new technologies and provide its technical staff with opportunities to improve their skill sets, this should be a conscious choice balanced against the other objectives of the organization.

Data-Driven Failures

After leadership-driven failures, interviewees identified data-driven failures as the second most common reason that AI projects end in failure. These difficulties manifested in a number of ways.

Many interviewees (30 of 50) discussed persistent issues with data quality. One interviewee noted,

80 percent of AI is the dirty work of data engineering. You need good people doing the dirty work—otherwise their mistakes poison the algorithms. The challenge is, how do we convince good people to do boring work?²⁸

Too Few Data Engineers

The lack of prestige associated with data engineering acts as an additional barrier: One interviewee referred to data engineers as “the plumbers of data science.”²⁹ Data engineers do the hard work of designing and maintaining the infrastructure that ingests, cleans, and transforms data into a format suitable for data scientists to train models on. Despite this, often the data scientists training the AI models are seen as doing “the real AI work,” while data engineering is looked down on as a menial task.³⁰ The goal for many data engineers is to grow their skills and transition into the role of data scientist; consequently, some organizations face high turnover rates in the data engineering group. Even worse, these individuals take all of their knowledge about the organization’s data and infrastructure when they leave. In organizations that lack effective documentation, the loss of a data engineer might mean that no one knows which datasets are reliable or how the meaning of a dataset might have shifted over time. Painstakingly rediscovering that knowledge increases the cost and time required to complete an AI project, which increases the likelihood that leadership will lose interest and abandon it.

Lack of Suitable Data

Additionally, in some cases, organizations lack the right kind of data to train AI models. This failure pattern is particularly common when the business is applying AI for the first time or to a new domain. Interviewees noted that business leaders often would be surprised to learn that their organization lacked sufficient data to train AI algorithms. As one interviewee put it, “They think they have great data because they get weekly sales reports, but they don’t realize the data they have currently may not meet its new purpose.”³¹ In many cases, legacy datasets were intended to preserve data for compliance or logging purposes. Unfortunately, structuring data for analysis can be quite different: It often requires considerable context about why things happened as opposed to simply what happened. For example, an e-commerce website might have logged what links users click on—but not a full list of what items appeared on the screen when the user selected one

or what search query led the user to see that item in the first place. This may mean that different fields need to be preserved, or different levels of granularity and quality may be necessary. Thus, even if an organization has a large quantity of historical data, that data may not be sufficient to train an effective AI algorithm.

Unbalanced Data

A related problem occurs when organizations have large quantities of data, but the data are unbalanced. For example, in health care applications, datasets may contain a large number of instances where a medical test correctly confirmed the absence of a rare cancer but only a handful of cases where the cancer was actually present. These conditions raise the risk of overfitting the data: The algorithm might excessively correlate the detection of these rare conditions with random, unrelated data characteristics from the handful of known cases. Gathering enough data to detect rare real-world events requires time, money, and patience.

Lack of Domain Understanding

Finally, several interviewees (10 of 50) noted that their lack of domain understanding could cause the failure of AI projects. Data scientists are rarely experts in the topics for which they are building their models: They require the assistance of subject-matter experts who can explain what the elements in the dataset mean and which ones are—and are not—important or might be unreliable. For example, a particular data field might appear at first glance to be highly relevant for training the AI model, but the data might be unreliable because they were manually entered by users who had little incentive to ensure that the data were of high quality. Unfortunately, in some cases, the subject-matter experts who are needed to support the AI team put up passive resistance to AI projects because they believe that these projects are intended to replace their jobs. In any case, without a detailed understanding of what the organization's data mean and which pieces of data are reliable and important, AI projects will often struggle to achieve the organization's aspirations for them.

Failures Due to Underinvestment in Infrastructure

One contributing factor to the numerous difficulties that organizations face in making their data ready for AI is the lack of investment in supporting infrastructure. Data engineering professionals need time to build up pipelines that can automatically clean data and continuously deliver fresh data to deployed AI models. Infrastructure investments ensure that these pipelines are automatically monitored to determine whether a data source changes formats or fails to arrive promptly. Organizations that quickly move from prototype to prototype often find that they are completely blind to failures that arise after the AI model has been completed and deployed. Robust infrastructure allows the engineering team to detect when a deployed model needs maintenance, which deployed models most urgently need maintenance, and what kind of maintenance action is required for each.

Additionally, investments in operations infrastructure ensure that AI models can be more quickly and easily deployed to production. Interviewees recommended investing in hiring ML engineers who have the specialized skills to build this infrastructure and speed up model deployments. Some interviewees noted that they had observed cases where AI models could not be deployed from test environments to production environments because the production environments were incompatible with the requirements of the model. In other cases, interviewees noted significant delays in deploying their completed models to end users because of a lack of robust infrastructure to automate the deployments. Ultimately, developing effective AI products requires more than just a data science team. Investing in data engineers and ML engineers can substantially shorten the time required to develop a new AI model and deploy it to a production environment, where it can actually help end users.

Failures Due to Immature Technology

Finally, interviewees observed that, in some cases, AI projects fail because some problems are still too difficult for AI algorithms to solve. The frequency of this

type of failure varies significantly depending on the type of use case for AI. For example, AI models are quite effective for many e-commerce or advertising use cases, but some intended applications for computer vision resist even the most rigorous and well-funded attempts to apply AI. One interviewee stated that AI algorithms are poorly suited to automating the internal processes of an organization—especially when subjective human judgment is required to determine how those processes should function. Leaders of an organization need to recognize that AI is not a magic tool that can fully automate any process or solve any problem. Some business use cases are a better fit for AI than others; understanding which problems are a good fit for AI and which are at or beyond the current state of the art can help organizations avoid costly and embarrassing failures.

Two Special Cases: Compute Power and Availability of Talent

Alongside data, talent and compute power are key prerequisites for the training of AI algorithms. No organization can expect to succeed in developing AI products without a strong foundation in each of these components. Unlike the situation for data, relatively few interviewees identified issues with the availability of either talent or compute power as the most frequent or impactful factors behind the failure of AI projects. However, because of the importance of these key inputs, we specifically asked the interviewees to discuss their perception of whether shortages in either of these areas contribute to the failure of AI projects.

Compute Power

Nearly all of the interviewees stated that compute power was not a limiting factor in their work. Most interviewees said that cloud computing providers offer substantial amounts of compute power for purchase on demand. Consequently, as long as the organization had adequately budgeted for the purchase of compute power, this was not a limiting factor in the development of AI algorithms. However, interviewees noted two exceptions to this rule. First, in some cases, companies think that their data

are too sensitive to upload to a cloud environment. This is particularly true in heavily regulated industries, such as finance or health care. However, even in these industries, many companies have successfully migrated their operations to the cloud in a way that preserves the security of their data. The second exception occurs when companies are operating at the edge of AI research. These are mostly large technology companies that are attempting to train their own LLMs. Even AI researchers who were not working on LLMs found that compute power might be rationed within the organization and that, in some cases, this would delay their ability to train or test models for a few days. However, several of these interviewees (4 of 50) expressed the belief that this would prove to be a temporary problem as graphics processing unit manufacturers ramp up production of their products.

Availability of Artificial Intelligence Talent

In contrast to the findings on compute power, when asked, many interviewees expressed the belief that the availability of AI talent does inhibit their work to some extent.³² Many interviewees noted that the overall availability of talent has improved in recent years as new master's programs in data science and bootcamps have produced graduates trained in the basic skills required to train AI algorithms. However, interviewees often noted that finding quality talent

Understanding which problems are a good fit for AI and which are at or beyond the current state of the art can help organizations avoid costly and embarrassing failures.

remains difficult. Many educational programs focus primarily on development of AI models as opposed to related skills in how to clean data, identify poor data, or deploy AI models to production environments. Consequently, interviewees said that they find it difficult to determine which recent graduates would be effective in a less pristine workplace environment where data might be dirty, undocumented, or unavailable.

Interviewees also observed that many companies want to hire AI workers with exposure to the latest techniques and models, even though relatively few of these companies truly need workers with these skills. Several interviewees (8 of 50) found that their organizations were most successful at hiring AI talent when they were prepared to identify talent with the potential to grow into the job as opposed to only hiring perceived “rockstars.”

Additionally, some interviewees noted the lack of consistency in industry titles as a barrier to hiring. The role of data scientist, in particular, can have radically different expectations and responsibilities across organizations. Open communication about exactly how the workplace functions is essential to ensure a good fit with potential new employees.

Agile Software Development and Artificial Intelligence

Finally, several interviewees (10 of 50) expressed the belief that rigid interpretations of agile software development processes are a poor fit for AI projects.³³ While the agile software movement never intended to develop rigid processes—one of its primary tenets is that individuals and interactions are much more important than processes and tools³⁴—many organizations require their engineering teams to universally follow the same agile processes. One interviewee noted that, in his experience, work items repeatedly had to either be reopened in the following sprint or made ridiculously small and meaningless to fit into a one-week or two-week sprint.³⁵ In particular, AI projects require an initial phase of data exploration and experimentation with an unpredictable duration. Interviewees recommended that instead of adopting established software engineering processes—which

often amount to nothing more than fancy to-do lists—the technical team should communicate frequently with their business partners about the state of the project. As one interviewee put it:

Stakeholders want to be a part of the process. They don't like it when you say, “it's taking longer than expected; I'll get back to you in two weeks.” They are curious.³⁶

Open communication builds trust between the business stakeholders and the technical team and increases the likelihood that the project will ultimately be successful.

Industry Interview Takeaways

As Table 3 illustrates, our industry interviews highlighted five leading root causes resulting in the failure of AI projects. First, business stakeholders often misunderstand—or miscommunicate—what problem needs to be solved using AI. Too often, organizations deploy trained AI models only to discover that the models have optimized the wrong metrics or do not fit into the overall workflow and context. Second, often the organization lacks the necessary data to adequately train an effective AI model. Third, in some cases, AI projects fail because they focus more on using the latest and greatest technology than on solving real problems for their intended users. Fourth, organizations often do not have adequate infrastructure to manage their data and deploy completed AI models, which increases the likelihood of project failure. Finally, in some cases, AI projects fail because the technology is applied to problems that are too difficult for AI to solve. AI is not a magic wand that can make any challenging problem disappear; in some cases, even the most advanced AI models cannot automate away a difficult task. These five root causes stood out in the industry interviews as the most common and most impactful reasons that data science teams in industry perceive AI projects as failing.

TABLE 3

Key Root Causes Highlighted in Industry Semistructured Interviews

Root Cause	Description
Leadership-driven failures	Leaders fail to communicate to the engineering team what problem they want to be solved and what metrics they need to optimize to solve it. Additionally, many leaders change priorities too rapidly to allow the engineering team to deliver effective AI models.
Data-driven failures	Organizations often lack sufficient high-quality data to train performant AI models. Leaders may not be prepared for the time and expense required to gather enough data to train an effective AI model.
Bottom-up–driven failures	Data scientists sometimes focus on using the most-advanced technology instead of finding the most effective way to solve the business problem.
Underinvestment in infrastructure	Inadequate infrastructure can lead to lower-quality data and longer deployment times for completed models. Underinvesting in infrastructure increases the risk that an AI project will fail.
Immature technology	In some cases, organizations attempt to apply AI to business problems that are beyond the state of the art for the technology.

Results of Interviews with Representatives of Academia

The academic AI research environment is different from the business environment. Academic research often focuses on developing new techniques through an integrative experimental process. *Failure* is hard to measure academically, as investigations into new computer algorithms or ML techniques are grounded in highly uncertain research areas. Unsurprisingly, we did not get a clear consensus on what AI failure is in academic research. During the interviews, we identified some root causes that might influence how academic researchers view AI project failures. A plurality of the interviews mentioned activity prestige (defined in the next section), data structures, and publication incentives as trends that would affect AI research. Additionally, we found that computing resources were not a large concern within the academic setting. We attribute this to the use of smaller datasets, more-efficient algorithms, and regular access to large computing, which are common in academia.

Activity Prestige

Given the demand for new AI projects, interviewees reported prioritizing projects that grab headlines, improve reputation, and increase institutional prestige. We refer to this as *activity prestige*, which is the amount of positive attention given to some

projects based on the public demand for those projects' outcomes. The interviewees indicated that higher-prestige projects often take priority over less attention-driven areas—even when the researcher believes that these other projects would be more useful or valuable (consequently making the project a failure from their perspective). This is not to say that the researchers personally believed the lower-prestige projects held less value; rather, there was an opportunity cost from focusing only on high-priority projects. In this context, a project was deemed successful if it resulted in prestigious outcomes.

This research builds on some key applied elements of information theory. This classic foundational research may have market implications. Experience is important in determining which activities researchers perceive as prestigious among their peers. Newer researchers are more focused on completing tenure-track requirements; consequently, they described feeling pressure to undertake AI projects that would result in publications. In contrast, recently tenured researchers often prioritized securing new or expanded funding sources. Finally, the well-established researchers (tenured for more than five years) we interviewed emphasized impact on new research lines as a success driver. From these observations, we found that while some researchers found publications to be a motivator, the focus on output-driven research overshadowed promising research areas that are more complicated and less linked to publishable outputs, making it less

likely that new researchers would presume possible innovation because these innovations would not be linked to tenure-track progress or valued within an academic institution. While these findings are anecdotal, future research could include nonacademic research organizations, such as think tanks or government-funded research affiliates, to see whether the publication pressure was constant in shaping the research areas.

Researchers also face incentives to undertake work that is more likely to result in boosts to their prestige. This means that they weigh the risk that a project will not result in a publication or additional funding when considering potential projects to undertake. Younger researchers also take into account the expected duration and then choose projects that have a shorter time horizon so that they will accumulate a more impressive record by the time they apply for tenure. Computer science research often focuses on improving technology, with little consideration of the practical application. At a company, AI production is more about practically applying the technology to business problems. This creates an incentive gap whereby academic researchers, especially newer ones, are more rewarded for taking on projects that have increased publication chances rather than real-world benefits. Once tenured, researchers have more freedom to take on riskier and longer-term projects that could have a greater impact.

Improper Data Structures

Improper data structures are related to data, bias, and collection. Data scale and distribution were marked as issues that could lead to a well-specified AI model. The most-prominent technical challenges were often in collecting and organizing the data needed to test a set of theoretical hypotheses that the researchers were seeking to better understand. Furthermore, sometimes data collections were biased or poorly constructed, as evidenced by many researchers noting that they went out of their way to collect the best possible samples. Sometimes, if data collection was infeasible, they would try to use synthetic or simulated datasets rather than using a biased data

collection. This point was especially emphasized by those doing biomedical research.

In research, there is a higher level of emphasis applied to diagnostics, performance, and measurement to highlight outcomes because the goal is to improve knowledge of computer science, which is distinctly different than applying that knowledge to business challenges. Academia's focus on the science of computers requires more-stringent procedures to ensure that the data are collected and reported to minimize harm to the users, often going through an internal review board and, if applicable, oversight controls. An example might be the application of an ML tool to patients with cancer, where the AI program will learn how to diagnose dark spots on the skin to predict the probability that they are cancerous. Patient data cannot be collected or stored within an openly accessed unstructured central repository but instead must be stored subject to patient record laws. For the enterprise of science, the quality of the data is paramount to grounding a theory or developing a new field.

Publication Incentives

Pressure to publish was repeatedly mentioned as a potential contributor to AI project failure. Many interviewees noted that when seeking tenure or building a research agenda, publication equals success. If an AI project did not result in a publication, then the project was not perceived as a success. It should be noted that a publication in this case could be any outward-facing engagement, such as a talk, conference paper, or proceeding. Given the high level of attention being placed on AI, there is an enormous demand for new ideas, concepts, and techniques, which further increases the institutional demand already placed on the researchers.

Some of the non-tenure-track respondents suggested that even if their AI projects were technically successful, they might not be successful in terms of advancing the researchers' academic career opportunities (making the project a failure from their perspective). Frequently, failure to publish on a project results from the discovery of a new and unforeseen technical problem, even though the act of problem identification itself is a contribution that might lead

to new insights or open up new avenues of research. Sometimes, the only way to identify a problem is to experiment and explore. However, even if a technical problem leads to a more promising research agenda, the interviewees noted that the project would still be considered a failure unless it resulted in an immediate publication, such as a conference proceeding or paper.

Other Findings

Almost none of the scholars mentioned the issues of access to computing resources, data storage, or a skilled labor force as limiting factors in their work. For access to computing, universities may have some of the largest, most powerful computers available or are dealing with smaller-scale datasets. Data storage was most often purchased via a private company and therefore often used a secure cloud service. The need for skilled labor in academic settings is easily managed through the academic institution’s focus on training and managing graduate and undergraduate students rather than seeking new sources of labor. Overall, although no respondent discussed quantity of available personnel being an issue, some noted that because laboratory workers were at different stages in their education, quality control was a key focus.

Respondents varied greatly in the different types of exogenous factors they identified. Several discussed

the recent surge in interest in AI technology, which has driven opportunities for more research but has also increased the demand for teachers and course offerings. Furthermore, at least one interviewee expressed the belief that the increasing popularity of LLMs could, in the long run, crowd out other types of AI research. However, other respondents noted that the interest in LLMs was complementary to other types of AI research. Many of the graduate students reported being optimistic about future AI research and feeling that it was an exciting time to be in the field.

Academic Interview Takeaways

Researcher participants noted that prestige, funding, and publication incentives played a large role in determining the success of an AI project within the academic space, as outlined in Table 4. While technical problems persist, they are often overcome with access to university-sponsored resources, including graduate student labor, computing services, and new hardware. When technical challenges arise, they often come from errors that are fixed during the research process or become new lines of research. In our interviews, we learned that when AI projects fail, they do so because of a misalignment in incentives rather than an overall technical barrier to product delivery. In short, overcoming an AI failure is more about humans than the machines.

TABLE 4
Key Root Causes Highlighted in Academic Semistructured Interviews

Root Cause	Description
Activity prestige	Researchers face pressure to do work that will be perceived by their peers as prestigious as opposed to work that could be impactful.
Improper data structures	Academic data often are either older or not collected with AI activities in mind; thus, the researchers will have issues with training data or insufficient amounts or quality of testing data.
Publication incentives	Researchers identified a project as a failure if it did not result in a publication, conference proceeding, or communication item, even if the project led to new AI research.

Industry Recommendations

Although AI projects can be challenging for any organization, failure is not inevitable. Leaders who want to avoid the most common mistakes cited by our AI experts should consider the following five recommendations that may help lead to successful AI implementation.

Ensure That Technical Staffs Understand Project Purpose and Domain Context

Misunderstandings and miscommunications about the intent and purpose of the project cause more AI projects to fail than any other factor. Both the business leaders and the engineers have a role to play in avoiding this outcome. Business leaders need to help their technical staff understand what they truly need the AI project to achieve and how the completed AI product will ultimately be used. They cannot assume that the engineering team can independently discover which design choices will make their product useful within its business context. At the same time, AI researchers and engineers need to earn the trust of their business stakeholders by keeping them apprised of their progress and project status, as well as any interim discoveries. Business leaders are often just as excited about the potential of AI as the engineers are, if not more so; appropriately including them in the journey helps ensure a successful outcome.

Before they begin any AI project, leaders should be prepared to commit each product team to solving a specific problem for at least a year.

Organizations should rethink the processes that they have in place to facilitate these connections and interactions among the various team members and stakeholders. Rigid interpretations of existing software development processes rarely suit the cadence of an AI project. Instead of forcing project teams to follow a uniform set of procedures designed for a different type of engineering, organizations should empower their teams to adapt their processes to fit their workloads. Ultimately, organizations will need to rediscover how to make the agile software development process be adaptive and—truly—agile.

Choose Enduring Problems

AI projects require time and patience to be completed successfully. Data scientists and data engineers need space to explore, understand, and curate the available data before attempting to train an AI model that will learn how to behave from those data. Rapidly shifting the team's priorities and chasing after the crisis or opportunity of the moment can lead to a string of AI projects being abandoned before they have a chance to deliver tangible results. Before they begin any AI project, leaders should be prepared to commit each product team to solving a specific problem for at least a year. If an AI project is not worthy of such a long-term commitment, it most likely is not worth committing to at all—especially because an AI project with an overly accelerated timeline is likely to fail without ever achieving its intended goal.

Focus on the Problem, not the Technology

Experienced engineers told us that successful project teams kept a clear focus on the business problem to be solved instead of the technology that would be used to solve it. Chasing the latest and greatest advances in AI for their own sake is one of the most frequent pathways to failure. Instead, an organization's leaders need to collaborate with the technologists to ensure that they select AI projects that are both a good fit for the technology and that solve a real problem for their intended user. No matter how impressive a new technology may appear, ultimately any technology—even AI—is simply a tool to be wielded rather than an end in and of itself.

Invest in Infrastructure

Data-related problems are among the top reasons AI projects fail. Building up data infrastructure to reliably clean, ingest, and monitor data streams can substantially improve an organization's data and ensure that more of its AI projects ultimately succeed. Additionally, investments in infrastructure to automatically deploy AI models allow organizations to deploy these models to production more rapidly and reliably, where they can deliver real benefits to real users. Too many businesses fail to recognize the value that these kinds of investments can provide; instead, they rapidly switch from one AI project to another without taking the time to invest in common tools that would make their data science teams more productive. Leaders often justify this strategy because technology and their businesses are changing too rapidly to make these kinds of investments. In reality, delaying investments in infrastructure makes AI projects take longer to complete and fail more often.

Understand Artificial Intelligence's Limitations

Finally, despite all the hype around AI as a technology, AI still has technical limitations that cannot always be overcome. Leaders cannot treat AI as a magic wand that can solve any problem or automate any process. Instead, leaders need to collaborate with their technical experts to choose projects that are a good fit for AI's capabilities and would deliver meaningful value to the organization. Leaders do not necessarily need to have a deep technical understanding of AI themselves, but they need to employ staff with a strong data science background when selecting objectives for their AI product teams. Simply assuming that AI can solve any problem risks setting the team up for failure.

Academic Recommendations

Overcome Data Collection Barriers Through Partnerships with Government

As we found in our interviews, academic researchers face some challenges collecting sufficient quantities of data to train effective AI models compared with their colleagues in industry. Academic use cases

require rigorous data-collection standards that often limit the amount of data that is available for research or may subtly bias the distribution of data samples collected. For example, collecting data for biometric tracking studies can require years of effort and still yield a relatively small dataset. In contrast, researchers in industry routinely collect and analyze much larger datasets. Consequently, academics could particularly benefit from large-scale datasets with well-documented collection procedures.

Local, state, and federal government agencies could play an important role in providing these foundational training sets as a public good. Government datasets meet academia's requirement that data must be collected in ways that meet the legal and ethical standards appropriate for academic research. Additionally, many government agencies collect data at the scale required to train sophisticated AI algorithms. In return, collaborating with academic researchers could help government agencies address their critical shortages of technical and AI talent. Such initiatives as Data.gov should be expanded and better funded to take advantage of these opportunities.

Expand Doctoral Programs in Data Science for Practitioners

In higher education, different roles have different incentives. Society benefits when academics are able to focus on basic research focusing on long-term problems. However, newer academics might need encouragement to tackle these challenging issues because they face short-term pressures to publish papers as quickly and reliably as possible. More-established academics, who have more job security, have greater freedom to take academic risks that may not pay off for several years. Governments and corporations should consider funding fellowships to enable innovative younger researchers to pursue longer-term research projects and free them from the pressure to constantly publish. Computer science and data science programs could learn from other disciplines, such as international relations and security studies, where practitioner doctoral programs often exist side by side at even the top-ranked universities to provide pathways for the most-advanced researchers to apply their findings to contemporary problems.

APPENDIX A

Industry Interview Template

This appendix includes the discussion questions used for industry interviewees.

1. Please introduce yourself. Please tell us a bit about your background and your experience with AI.
2. In your experience, when an AI project fails or encounters significant difficulties, what are the most common root causes for those problems?
3. What kinds of technical causes can contribute to the failure of an AI/ML project?
 - Problems with data
 - Using the wrong approach
 - Other technical root causes?
4. What kinds of nontechnical causes can contribute to the failure of an AI/ML project?
 - Direction from leadership
 - Internal politics/bureaucracy
 - Other nontechnical root causes
5. Have any of these areas ever caused problems for your AI/ML projects?
 - Availability of talent (unable to hire the right people)
 - Availability of compute
6. Are there differences in bottom-up reasons why AI/ML projects fail compared to top-down reasons why AI/ML projects fail?
7. What kinds of problems have you encountered as AI/ML teams expand from a small core group to a larger team?
8. Does the rapidly changing nature of AI/ML itself cause problems for your AI/ML projects?
9. What else should we have asked you?

APPENDIX B

Academia Interview Template

This appendix includes the discussion questions used for academic interviewees.

1. Please introduce yourself. Please tell us a bit about your background and some key elements that establish your knowledge base associated with AI technologies.
 - **Probe 1.1:** How many projects involving AI have you worked on? What did those projects aim to accomplish?
2. Tell us a story about your most relevant experience where an AI project was not successful. We are interested in your experience in why an unsuccessful project has challenges or roadblocks and ultimately was not able to move to completion.
 - **Probe 2.1:** What do you think were the root causes of that failure? Are there any elements that stand out as critical inflection points that put stress on the project? What, if any, were the results of these inflection points?
 - **Probe 2.2:** What are the first-, second-, and third-order effects of the key failure points? Are there clear patterns within the failures that could be observed from your vantage point?
3. Consider the difficulty of developing AI systems. We would like to explore root technical causes of AI failure. The complexity of AI development cycles and the novelty of its application are often a source of reported technical difficulty.
 - **Probe 3.1:** How did data structure, collection, and capture approaches affect the failure of developing an AI system? Was this a data problem?
 - **Probe 3.2:** Given the complexity of the modeling, was modeling choice or specification of selection result a significant contributing factor to an AI failure? Was this a model or development problem?
 - **Probe 3.3:** Do the resulting workflow, development processes, and shifts in technology platforms present a challenge that significantly contributes to the likelihood of AI failure? Do changes in the technology standards matter?
4. We now want to pivot to examining the organizational roots of AI failures within an organizational structure and the human, labor, and capital dynamics that might shape AI failure.
 - **Probe 4.1:** What workforce challenges caused difficulties for your projects?
5. Exogenous and systematic factors may shape if an AI is more or less likely to fail. How, in your experience, do network externalities, labor or capital shortages, market failures, market forces, and/or regulatory frameworks influence AI failure?
6. In your view, what are the most common reasons that you have observed or experienced that resulted in AI failure?
7. Thank you for coming today and for sharing your opinions with us. We hope you enjoyed the discussion today.

Notes

¹ For this project, we focused on the machine learning (ML) branch of AI because that is the technology underpinning most business applications of AI today. This includes AI models trained using supervised learning, unsupervised learning, or reinforcement learning approaches and large language models (LLMs). Projects that simply used pretrained LLMs (sometimes known as *prompt engineering*) but did not attempt to train or customize their own were not included in the scope of this work.

² Maslej et al., *Artificial Intelligence Index Report 2023*. The report notes that private-sector AI investment decreased in 2022 for the first time in a decade.

³ *Midsized corporations* are those with between \$100 million and \$499 million in annual revenue.

⁴ Patience and Immerman, *2023 Global Trends in AI Report*.

⁵ Vincent, “Why the Pentagon Didn’t Request Higher Funding for AI in Fiscal 2025.”

⁶ Anagnostopoulos et al., “How Artificial Intelligence Can Power Clinical Development.”

⁷ Musani, “Decking the Aisles with Data: How Walmart’s AI-Powered Inventory System Brightens the Holidays.”

⁸ Copp, “An AI-Controlled Fighter Jet Took the Air Force Leader for a Historic Ride.”

⁹ Seffers, “Undersea Combat Includes Way More Than Submarines.”

¹⁰ Clark, “Pentagon Official Lays Out DoD Vision for AI.”

¹¹ Cisco AI Readiness Index.

¹² McKendrick, “Executives and Managers Welcome AI, but Struggle with It.”

¹³ Kahn, “Want Your Company’s AI Project to Succeed?”

¹⁴ Bojinov, “Keep Your AI Projects on Track.”

¹⁵ Budgen, *Software Design*; Ambler, *Process Patterns*.

¹⁶ Sjödin et al., “How AI Capabilities Enable Business Model Innovation.”

¹⁷ Standing et al., “The Attribution of Success and Failure in IT Projects.”

¹⁸ Westenberger, Schuler, and Schlegel, “Failure of AI Projects.”

¹⁹ Bojinov, “Keep Your AI Projects on Track”; Schmelzer, “The One Practice That Is Separating the AI Successes from the Failures.”

²⁰ Mittal, Saif, and Ammanath, *Deloitte’s State of AI in the Enterprise, 5th Edition Report*.

²¹ Chui et al., “The State of AI in 2022—and a Half Decade in Review.”

²² Schlegel, Schuler, and Westenberger, “Failure Factors of AI Projects.”

²³ *Individual contributor* is a common term used in industry to denote an employee who is not responsible for managing other employees. Individual contributors can have various levels of seniority and work experience. For this project, the term *manager* includes all levels of management, including director, vice president, and other titles. It includes both managers of individual contributors and managers of managers.

²⁴ Many interviewees had worked for more than one company during their career.

²⁵ The majority of the rest of our candidate pool were nonresponders. Approximately two dozen individuals responded that they were interested in participating in the study but did not set up an interview time before we ended data collection.

²⁶ *Meaningful number* here means more than one-quarter but less than one-half of the interviewees identified these root causes as among the primary reasons for AI project failures.

²⁷ Study participant, interview with the authors, October 6, 2023.

²⁸ Study participant, interview with the authors, September 12, 2023.

²⁹ Study participant, interview with the authors, September 28, 2023.

³⁰ Study participant, interview with the authors, November 20, 2023.

³¹ Study participant, interview with the authors, October 10, 2023.

³² Seven of 50 interviewees stated that talent availability was a difficulty in AI projects. An additional 19 of 50 stated that while talent availability was not a problem overall, the lack of high-quality talent was a limitation in their AI projects.

³³ *Agile software development* is defined by the U.S. Government Accountability Office as “an approach to software development in which software is built incrementally and is continuously evaluated on functionality, quality, and customer satisfaction” (U.S. Government Accountability Office, *Science & Tech Spotlight: Agile Software Development*).

³⁴ Beck et al., “Manifesto for Agile Software Development.”

³⁵ In agile software development, *sprints* are a short, time-boxed period in which the engineering team commits to completing a specific set of work. Sprints typically last one to four weeks.

³⁶ Study participant, interview with the authors, October 6, 2023.

References

- Ambler, Scott W., *Process Patterns: Building Large-Scale Systems Using Object Technology*, Cambridge University Press, 1998.
- Anagnostopoulos, Chris, David Champagne, Thomas Devenyns, Alex Devereson, and Heikki Tarkkila, "How Artificial Intelligence Can Power Clinical Development," McKinsey & Company, November 22, 2023.
- Beck, Kent, Mike Beedle, Arie van Bennekum, Alistair Cockburn, Ward Cunningham, Martin Fowler, James Grenning, Jim Highsmith, Andrew Hunt, Ron Jeffries, et al., "Manifesto for Agile Software Development," Agile Alliance, February 2001.
- Bojinov, Iavor, "Keep Your AI Projects on Track," *Harvard Business Review*, November–December 2023.
- Budgen, David, *Software Design*, 2nd ed., Addison-Wesley, 2003.
- Chui, Michael, Bryce Hall, Helen Mayhew, Alex Singla, and Alex Sukharevsky, "The State of AI in 2022—and a Half Decade in Review," McKinsey & Company, December 6, 2022.
- Cisco AI Readiness Index, homepage, undated. As of June 19, 2024:
https://www.cisco.com/c/m/en_us/solutions/ai/readiness-index.html
- Clark, Joseph, "Pentagon Official Lays Out DoD Vision for AI," *DoD News*, February 21, 2024.
- Copp, Tara, "An AI-Controlled Fighter Jet Took the Air Force Leader for a Historic Ride. What That Means for War," Associated Press, May 3, 2024.
- Kahn, Jeremy, "Want Your Company's A.I. Project to Succeed? Don't Hand It to the Data Scientists, Says This CEO," *Fortune*, July 26, 2022.
- Maslej, Nestor, Loredana Fattorini, Erik Brynjolfsson, John Etchemendy, Katrina Ligett, Terah Lyons, James Manyika, Helen Ngo, Juan Carlos Niebles, Vanessa Parli, Yoav Shoham, Russell Wald, Jack Clark, and Raymond Perrault, *Artificial Intelligence Index Report 2023*, AI Index Steering Committee, Institute for Human-Centered AI, Stanford University, 2023.
- McKendrick, Joe, "Executives and Managers Welcome AI, but Struggle with It," *Forbes*, August 24, 2023.
- Mittal, Nitin, Irfan Saif, and Beena Ammanath, *Deloitte's State of AI in the Enterprise, 5th Edition Report*, Deloitte, October 2022.
- Musani, Parvez, "Decking the Aisles with Data: How Walmart's AI-Powered Inventory System Brightens the Holidays," Walmart Global Tech, October 25, 2023.
- Patience, Nick, and David Immerman, *2023 Global Trends in AI Report*, S&P Global, August 2023.
- Schlegel, Dennis, Kajetan Schuler, and Jens Westenberger, "Failure Factors of AI Projects: Results from Expert Interviews," *International Journal of Information Systems and Project Management*, Vol. 11, No. 3, 2023.
- Schmelzer, Ron, "The One Practice That Is Separating the AI Successes from the Failures," *Forbes*, August 14, 2023.
- Seffers, George, "Undersea Combat Includes Way More Than Submarines," *Signal Magazine*, February 1, 2022.
- Sjodin, David, Vinit Parida, Maximilian Palmié, and Joakim Wincent, "How AI Capabilities Enable Business Model Innovation: Scaling AI Through Co-Evolutionary Processes and Feedback Loops," *Journal of Business Research*, Vol. 134, September 1, 2021.
- Standing, Craig, Andrew Guilfoyle, Chad Lin, and Peter E. D. Love, "The Attribution of Success and Failure in IT Projects," *Industrial Management and Data Systems*, Vol. 106, No. 8, October 2006.
- U.S. Government Accountability Office, *Science & Tech Spotlight: Agile Software Development*, GAO-20-713SP, September 29, 2020.
- Vincent, Brandi, "Why the Pentagon Didn't Request Higher Funding for AI in Fiscal 2025," *DefenseScoop*, March 11, 2024.
- Westenberger, Jens, Kajetan Schuler, and Dennis Schlegel, "Failure of AI Projects: Understanding the Critical Factors," *Procedia Computer Science*, Vol. 196, 2022.



About This Report

Leaders across all sectors are increasingly recognizing the importance of artificial intelligence and machine learning (AI/ML) for their organization, believing or hoping that AI will help them achieve formerly impossible feats and rise above their peers. Yet all too often, AI projects flounder or never get off the ground. To investigate why AI projects fail, RAND researchers conducted an exploratory study, interviewing data scientists and engineers with at least five years of experience building AI/ML models in industry or academia and synthesized the experts' experiences to develop a set of recommendations for smart implementation of AI. The findings and recommendations of this report should be of interest to the U.S. Department of Defense, which has been actively looking for ways to use AI, along with other leaders in government and the private sector who are considering using AI/ML. The lessons from earlier efforts to build and apply AI/ML will be helpful for others to avoid the same pitfalls.

The research reported here was completed in April 2024 and underwent security review with the sponsor and the Defense Office of Prepublication and Security Review before public release.

RAND National Security Research Division

This research was conducted within the Acquisition and Technology Policy Program of the RAND National Security Research Division (NSRD), which operates the RAND National Defense Research Institute (NDRI), a federally funded research and development center (FFRDC) sponsored by the Office of the Secretary of Defense, the Joint Staff, the Unified Combatant Commands, the Navy, the Marine Corps, the defense agencies, and the defense intelligence enterprise. This research was made possible by NDRI exploratory research funding that was provided through the FFRDC contract and approved by NDRI's primary sponsor.

For more information on the RAND Acquisition and Technology Policy Program, see www.rand.org/nsrd/atp or contact the director (contact information is provided on the webpage).

Acknowledgments

We would like to thank Jeannette Tsuei for her assistance with this project and Christopher Mouton, Yun Kang, Caitlin Lee, and Megan McKernan for their support of the project. We would also like to thank our reviewers (Justin Lynch, Pete Schirmer, Michael Kennedy, Chad Ohlandt, Bonnie Triezenberg, Jeff Alstott, and Aaron Frank) for all of their helpful insights and questions. Finally, we would like to thank all of our interviewees for their time and their perspectives about this important topic.

RAND is a research organization that develops solutions to public policy challenges to help make communities throughout the world safer and more secure, healthier and more prosperous. RAND is nonprofit, nonpartisan, and committed to the public interest.

Research Integrity

Our mission to help improve policy and decisionmaking through research and analysis is enabled through our core values of quality and objectivity and our unwavering commitment to the highest level of integrity and ethical behavior. To help ensure our research and analysis are rigorous, objective, and nonpartisan, we subject our research publications to a robust and exacting quality-assurance process; avoid both the appearance and reality of financial and other conflicts of interest through staff training, project screening, and a policy of mandatory disclosure; and pursue transparency in our research engagements through our commitment to the open publication of our research findings and recommendations, disclosure of the source of funding of published research, and policies to ensure intellectual independence. For more information, visit www.rand.org/about/research-integrity.

RAND's publications do not necessarily reflect the opinions of its research clients and sponsors. **RAND**® is a registered trademark.

Limited Print and Electronic Distribution Rights

This publication and trademark(s) contained herein are protected by law. This representation of RAND intellectual property is provided for noncommercial use only. Unauthorized posting of this publication online is prohibited; linking directly to its webpage on rand.org is encouraged. Permission is required from RAND to reproduce, or reuse in another form, any of its research products for commercial purposes. For information on reprint and reuse permissions, please visit www.rand.org/pubs/permissions.

For more information on this publication, visit www.rand.org/t/RR-A2680-1.

© 2024 RAND Corporation

www.rand.org