MIS QUARTERLY EXECUTIVE

Successfully Mitigating AI Management Risks to Scale AI Globally

Many firms struggle to scale today's generative and predictive AI systems effectively because their machine learning-based working mechanisms amplify general technology management challenges and create entirely new ones. Based on an in-depth case study of industrial AI pioneer Siemens AG, we describe how to successfully mitigate five critical technology management risks to scale AI globally, and provide recommendations for creating company-wide business impacts with machine learning-based AI systems.[1,2]

Thomas Hutzschenreuter
Technical University of Munich, School of Management (Germany)

Tim Lämmermann
Technical University of Munich, School of Management (Germany)

Alexander Sake
Siemens AG (Germany)

Helmuth Ludwig
Southern Methodist University (U.S.)

The Need to Rethink Technology Management Practices for Today's AI

AI systems continue to expand the scope of business opportunities available to firms . Today, these systems are primarily exploited for tasks benefiting from content generation, prediction-making, or a combination of both.[3] Whereas generative AI functionalities focus on producing new content such as text, image, video, audio or code, predictive AI produces estimates of a future phenomenon that can be useful for classification, forecasting and decision-making tasks.[4] Because of these promising application possibilities, global AI investments are projected to exceed $600 billion in the coming years.[5]

1 Mary Lacity is the accepting senior editor for this article.
2 The authors would like to express their sincerest gratitude to Mary Lacity for her extremely valuable comments and ideas and her excellent developmental guidance throughout the article's review process. We also thank the review team for their very constructive comments and suggestions, which have significantly improved this article, and Siemens for its excellent collaboration during our research.
3 For insights on AI adoption areas, see McElheran, K., Li, J. F., Brynjolfsson, E., Kroff, Z., Dinlersoz, E., Foster, L. and Zolas, N. "AI Adoption in America: Who, What, and Where," *Journal of Economics & Management Strategy* (33:2), January 2024, pp. 375-415.
4 See Hillebrand, L., Raisch, S. and Schad, J. "Managing With Artificial Intelligence: An Integrative Framework," *Academy of Management Annals* (19:1), January 2025, pp. 343-375.
5 Shirer, M. *Worldwide Spending on Artificial Intelligence Forecast to Reach $632 Billion in 2028, According to a New IDC Spending Guide*, International Data Corporation, August 19, 2024, available at https://www.idc.com/getdoc.jsp?containerId=prUS52530724.

DOI: 10.17705/2msqe.00117

However, practitioners and scholars regularly emphasize how difficult it is to create measurable business impacts with today's AI systems.[6] Indeed, because of various complex technology management challenges to successfully implement and exploit intelligent machines, over 70% of AI implementation projects fail to make an impact.[7] Without a sufficient understanding of critical AI management risks and corresponding mitigation practices, firms cannot strategically scale AI-based technologies globally.[8]

To successfully deploy today's AI applications, firms must consider how their inherent working mechanisms differ from those of traditional information systems. Today's AI systems primarily rely on machine learning, which enables algorithms to autonomously learn to perform cognitively demanding content-generation and prediction-making tasks.[9] By probabilistically detecting correlations within task-related training data, machine learning-based AI systems can form powerful decision rules that might even surpass human knowledge. In contrast, traditional information technologies and earlier AI applications are often more deterministic because they primarily rely on manually programmed decision rules. Hence, they merely reflect human-inserted task knowledge but cannot learn independently.

6 Note that "early AI" applications, such as expert systems, did not typically apply machine learning methods but were entirely manually coded. This is why we prefer the term "today's" or "contemporary" AI systems to refer to machine learning-based applications (including both generative and predictive AI functionalities).

7 See, for example: 1) Ångström, R. C., Björn, M., Dahlander, L., Mähring, M. and Wallin, M. W. "Getting AI Implementation Right: Insights from a Global Survey," *California Management Review* (66:1), August 2023, pp. 5-22; 2) "The Widespread Adoption of AI by Companies Will Take a While," *The Economist*, March 29, 2023, available at https://www.economist.com/leaders/2023/06/29/the-widespread-adoption-of-ai-by-companies-will-take-a-while; and 3) *Why Do 87% of Data Science Projects Never Make It Into Production?*, VentureBeat, July 19, 2019, available at https://venturebeat.com/ai/why-do-87-of-data-science-projects-never-make-it-into-production/.

8 See: 1) Hutzschenreuter, T. and Lämmermann, T. "What Is Your AI Strategy? Systematically Integrating Self-Learning Technologies into Your Business Strategy," *Academy of Management Perspectives*, 2025, published online January 2025; and 2) Lacity, M. and Willcocks, L. P. "Becoming Strategic with Intelligent Automation," *MIS Quarterly Executive* (20:2), June 2021, pp. 169-182.

9 For more details, see: 1) Berente, N., Recker, J., Gu, B. and Santhanam, R. "Managing Artificial Intelligence," *MIS Quarterly* (45:3), September 2021, pp. 1433-1450; and 2) Jordan, M. I. and Mitchell, T. M. "Machine learning: Trends, perspectives, and prospects," *Science* (349:6245), July 2015, pp. 255-260.

As a result, the distinctive working mechanisms of machine learning-based AI can amplify general technology management risks[10] and create entirely new ones. Hence, to strategically scale AI,[11] firms must systematically combine existing technology management practices with those that explicitly address the inherent nature of machine learning. To provide practical recommendations on how to achieve this, we investigated the following two research questions:

1. What are the critical technology management risks for machine learning-based AI systems?
2. How can firms effectively mitigate these risks to scale AI globally?

To answer these questions, we conducted an in-depth case study of industrial AI[12] pioneer Siemens AG (our research method is described in the Appendix). In this article, we describe how Siemens successfully mitigates five crucial technology management risks to strategically scale generative and predictive AI applications globally. Based on Siemens's risk mitigation practices, we provide practical recommendations for creating company-wide business impacts with today's machine learning-based AI systems. Each of these recommendations can generally help firms in effectively managing machine learning projects.

Siemens and Its Industrial AI Journey

Siemens AG is a global technology leader with the strategic vision to shape technology landscapes worldwide.[13] By focusing on

10 Schiffer, S., Mocker, M. and Teubner, A. "Managing IT Challenges when Scaling Digital Innovations," *MIS Quarterly Executive* (22:3), September 2023, pp. 209-218.

11 See, for instance: 1) Sagodi, A., van Giffen, B., Schniertshauer, J., Niehues, K. and vom Brocke, J. "How Audi Scales Artificial Intelligence in Manufacturing," *MIS Quarterly Executive* (23:2), June 2024, pp. 185-204; and 2) van Giffen, B. and Ludwig, H. "How Siemens Democratized Artificial Intelligence," *MIS Quarterly Executive* (22:1), March 2023, pp. 1-21.

12 Industrial AI is the exploitation of generative and predictive AI functionalities in manufacturing contexts. See Peres, R. S., Jia, X., Lee, J., Sun, K., Colombo, A. W. and Barata, J. "Industrial Artificial Intelligence in Industry 4.0—Systematic Review, Challenges and Outlook," *IEEE Access* (8), December 2020, pp. 220121-220139.

13 See *Technology to Transform the Everyday for Everyone*, Siemens, available at https://www.siemens.com/global/en/company/about/strategy/technology-to-transform-the-everyday.html

Table 1: Brief Overview of Siemens AG

Headquarters	Munich, Germany
Business focus	Engineering high-quality technology systems and offering customer-centric industrial services
Customer target	Mainly business customers (B2B)
Industry segments	Automation, infrastructure, mobility and healthcare solutions
Industry trends	Digitalization, demographic change, urbanization, globalization/ localization, environmental change, resource efficiency
Geographic markets	Worldwide planning, production, administration and sales facilities
Employees (2024 fiscal year)	312,000
Revenue (2024 fiscal year)	€75.9 billion
Net income (2024 fiscal year)	€9.0 billion
Strategic priority for AI	The global scaling of generative and predictive AI is a crucial board topic for the firm
Strategic AI vision	Differentiating the firm from global competitors through outstanding industrial AI competencies
AI implementation focus	Internally: Systematically deploying AI in global value chain functions like R&D, production, technical project management, sales and marketing Market offerings: Enhancing existing industry products and services, but also developing new technology offerings to gain first-mover advantages
Level of AI skills	Strong technology management skills due to the firm's long-standing experience with machine learning along its AI journey

industrial automation, infrastructure, mobility, and healthcare products and services, Siemens accelerates customers' digital and sustainable transformations, making factories more efficient, cities more livable and transportation more ecofriendly. In its 2024 fiscal year[14] the firm's revenue was €75.9 billion (approximately $88.7 billion),[15] with a net income of €9.0 billion and 312,000 employees worldwide (Table 1 provides a brief overview of Siemens).

Siemens has been working with AI for over half a century and has developed strong industrial technology competencies over time.[16] For example, after the era of expert and decision-support systems during the 1970s and 1980s, the firm introduced machine learning-based predictive AI systems in the 2000s. These systems performed tasks like estimating revenue streams, improving logistic routes and interconnecting industrial shop floor machines. Since 2022, Siemens has increasingly explored strategic business opportunities enabled by generative AI systems, as emphasized by the CEO Dr. Roland Busch: "[Generative AI] has the potential to revolutionize the way companies design, develop, manufacture and operate."[17]

Throughout its AI journey, Siemens has focused on industrial AI applications that can significantly differentiate it from competitors. This is underlined by the company's chief technology officer (CTO), Dr. Peter Körte:

14 Siemens's fiscal year 2024 began in October 2023 and ended in September 2024.

15 Currency conversion as of July 2025.

16 Scheibenzuber, M. *Tracing the AI Family Tree*, Siemens, available at https://www.siemens.com/global/en/company/about/history/stories/tracing-the-ai-family-tree.html. More insights on Siemens's strategic AI perspective can be found in van Giffen, B. and Ludwig, H., op. cit., March 2023.

17 *Siemens and Microsoft to Work Together on AI Project*, Reuters, October 31, 2023, available at https://www.reuters.com/technology/siemens-microsoft-work-together-ai-project-2023-10-31/.

"Industrial AI is a game-changer that will create significant positive impact in the real world across all industries."[18] From a functional perspective, industrial AI can improve operational efficiencies in global R&D, production, project management, marketing and sales departments. From a product perspective, there is much potential for including AI features in existing core engineering products or services. Further opportunities are provided by the emergence of new AI-driven business models.

To date, Siemens has registered more than 3,700 AI patents, resulting in a large number of successfully implemented machine learning systems.[19] Moreover, in recent years, the company has identified more than 400 novel industrial AI use cases, the majority of which are related to generative AI. Given the multi-year nature of complex technology implementations, since 20% to 25% of use cases were selected for a pilot phase, Siemens currently has dozens of novel AI projects in pilot, deployment or scaling phases.[20] However, many AI projects have faced substantial challenges because of their different technology management risks. But even these projects are very important for the firm because they offer valuable learning. Indeed, the experiences gained from both project successes and failures have enabled Siemens to acquire outstanding industrial AI competencies.

Notably, Siemens prioritizes the ethical aspects for human and organizational system stakeholders when developing, implementing and scaling AI. For instance, decision-making algorithms must provide trustworthy and non-biased outputs, intelligent transportation systems must reliably protect passengers, and business analytics tools must handle personal information sensitively. Moreover, the company considers essential legal and quality assurance aspects to adequately guide current and future implementation projects. To ensure that these aspects are fully considered, Siemens established a company-wide AI governance framework, underlining the obligation to ensure algorithmic accountability, fairness, cybersecurity, understandability and privacy needs for implemented AI systems.[21]

Description of Three Globally Scaled Industrial AI Projects at Siemens

To illustrate Siemens's industrial AI leadership, we describe three projects (summarized in Table 2) that have been scaled company-wide. These three exemplary projects emphasize the firm's existing competencies in both generative and predictive AI, though it is important to stress that Siemens's AI journey is ongoing. The firm continuously searches for new use cases to further improve its market offerings or internal business processes.

Industrial Copilot

Following the increasing availability of generative AI applications like ChatGPT, Stable Diffusion and Co. over the last few years, Siemens started to discuss which deployment contexts could perceivably benefit from content-generation functionalities.[22] As a consequence, to counteract the increased global competition in engineering markets and mitigate potential labor shortages, the firm identified the opportunity to differentiate itself with a novel industry offering that enables engineers to program industrial shop floor machines with a generative AI assistant. The resulting Industrial Copilot functions as a direct communication interface between humans and programmable manufacturing machines based on natural language processing. Instead of complex

18 For additional insights on Siemens's industrial AI vision, see *Siemens Unveils Breakthrough Innovations in Industrial AI and Digital Twin Technology at CES 2025*, Siemens, January 6, 2025, available at https://newsroom.sw.siemens.com/en-US/siemens-ces-2025/.

19 These include both AI-based and AI-enhanced digital systems. For more information on some exemplary use cases, see *Discover Industrial-grade AI Use Cases Across Various Industries*, Siemens, available at https://www.siemens.com/global/en/products/automation/topic-areas/artificial-intelligence-in-industry.html#DiscoverindustrialgradeAIusecasesacrossvariousindustries.

20 These include both global lighthouse and local niche projects.

21 For more on AI governance at Siemens, see Beitinger, G. *Revolutionizing Manufacturing: Navigating the Artificial Intelligence Landscape for Efficiency, Ethics, and Growths*, Siemens, May 5, 2024, available at: https://blog.siemens.com/2024/05/revolutionizing-manufacturing-navigating-the-artificial-intelligence-landscape-for-efficiency-ethics-and-growth/. For further insights on the relevance of AI governance for firms, see Schulte-Derne, D. and Gnewuch, U. "Translating AI Ethics Principles into Practice to Support Robotic Process Automation Implementation," *MIS Quarterly Executive* (23:2), June 2024, pp. 187-203.

22 For more insights on Siemens's Industrial Copilot, see *Industrial Copilots: Generative AI-Powered Value Chain Optimization*, Siemens, available at https://www.siemens.com/global/en/products/automation/topic-areas/artificial-intelligence-in-industry/industrial-copilot.html.

Table 2: Three Exemplary Globally Scaled AI Projects at Siemens

AI Project	AI Functionalities	Use Case Description	Global Business Impact
Industrial Copilot	Generative AI	Used to (partially) automate machine coding tasks in production facilities. Via a natural language interface, robotic machine movements can be steered with text commands. This, in turn, reduces setup times, enhances production capacities and helps to find manual coding bugs.	The ultimate goal is to gain first-mover advantages with generative AI, as this innovative market offering can be sold to industrial customers worldwide. The Industrial Copilot can also publicly showcase Siemens's very advanced AI competencies.
Virtual Prototyping	Predictive AI	Supports R&D tasks with AI-powered virtual simulations for complex machine prototypes (immersive engineering). By creating digital twins, the tool significantly streamlines the comparison of potential design candidates. This is especially beneficial for complex prototypes, requiring extensive space and technological infrastructure for physical tests.	Potential to reduce resource needs for many different R&D tasks internally. Additionally, the tool can be offered as an innovative customer offering in industrial technology markets.
SiemensGPT	Combination of generative and predictive AI	Provides a company-wide platform that bundles various generative and predictive AI functionalities. Through a chatbot interface, employees can search for relevant project information, generate business reports or exchange information about potential process improvements.	Impacts the firm's technical project and process management skills. Specifically, because the platform is accessible company-wide, it enables employees to innovate by building intelligent work companions themselves.

coding languages, engineers can, to some extent, tell a machine what to do. For instance, customers can use the AI interface to better coordinate and align robotic machine movements involved in production tasks or initiate timely maintenance processes. Using the Industrial Copilot in these ways can significantly reduce machine setup or repair times, increase production capacities and help avoid programming code bugs.

Because of potential first-mover advantages, the Industrial Copilot received much management support from the beginning. Though it took some time to fine-tune[23] the pretrained large language models[24] with industry-specific machine data, the resulting customer product is being offered to more and more industry markets, thus further underlining Siemens's ability to innovate with industrial AI solutions.

23 Fine-tuning adapts the self-learned decision rules of a pre-developed machine learning model to a specific implementation context by feeding it with additional task-specific training data.

24 Large language models are machine learning algorithms trained on vast amounts of text data to understand, analyze and generate human-like language. For more information, see Feuerriegel, S., Hartmann, J., Janiesch, C. and Zschech, P. "Generative AI," *Business & Information Systems Engineering* (66:1), February 2023, pp. 111-126.

Virtual Prototyping

The Virtual Prototyping system supports industrial R&D tasks.[25] The underlying use case problem was that product simulations or experiments have to be conducted in physical environments. This required space and expensive infrastructure, particularly in the case of complex machine prototypes. To counteract this problem, Siemens recognized that AI functionalities could be used to create virtual machine simulations with, for example, digital twins of manufacturing robots, engines or medical devices, a process known as "immersive engineering." The AI system was trained on technical performance metrics such as material friction, deformation, heat release and durability to compare the characteristics of different design candidates in varying environmental conditions. Given its cost savings and product-quality benefits, Virtual Prototyping has become an internal standard in Siemens's R&D units globally, enabling them to decrease development times and better meet customer requirements.

SiemensGPT

The SiemensGPT project was initiated when technology experts noticed the potential of bundling several generative and predictive AI functionalities in a central company platform. This platform facilitates the project work of employees regardless of their job type. Users can easily interact with the system via a chatbot interface, removing the need to become familiar with complex technological background processes.

The SiemensGPT platform is connected to different tools, including a web search assistant, code interpreters, internal databases and Siemens's organizational knowledge base. Employees can, for instance, use different AI-based functionalities to search for relevant industry project information, summarize and generate technical documents, or mutually brainstorm about potential process improvements. At this point, SiemensGPT has access to over 90,000 internal documents and runs on dozens of large language models from different providers. Because it has been made accessible to everyone in the global organization, it already has more than 70,000 active users worldwide.

Critical Technology Management Risks in Siemens's Industrial AI Journey

Throughout its industrial AI journey, Siemens encountered five different but equally critical categories of technology management risks associated with developing, implementing and exploiting AI systems. Below, we describe each risk category in detail and highlight its criticality for Siemens's global AI projects.

Risk 1: Missing or Falsely Evaluated Potential AI Opportunities

The first critical risk is related to the identification of promising AI use cases. Because the field of AI is complex and dynamically evolving, there is always the risk of not identifying valuable AI adoption ideas for both internal processes and customer products. Though novel use cases are the driving force behind Siemens's AI journey, overlooking opportunities could invite competitors to capitalize on them, potentially putting Siemens at a competitive disadvantage. To avoid this, the firm needs to be aware of novel technological opportunities, industry innovations, regulations, internal suggestions and potential service providers. Collecting this information, however, is very demanding and time-consuming: "Even [as] … AI experts, we are … [sometimes] surprised about the new things and possibilities. … We are just trying to … find [the best] ways of using predictive or generative AI" (Principal AI technology expert).

For example, to create the Industrial Copilot, Siemens had to collect information about generative AI's programming abilities, target machinery, available hardware components, estimated market sizes, potential competitor offerings, etc. This included accessing internal information sources from the R&D and marketing departments, as well as external sources such as hardware and software providers, market research firms and press releases.

A related hurdle is the appropriate evaluation of identified AI opportunities because ineffective

25 For more details on *Virtual Prototyping*, see Advanced Prototyping, Siemens, available at https://xcelerator.siemens.com/global/en/all-offerings/services/a/advanced-prototyping.html.

use cases typically lead to investment losses and technology frustration. Therefore, it is necessary to apply very different use case evaluation criteria. The evaluations also include make-or-buy decisions because the firm must determine whether to carry out AI endeavors internally or seek assistance from external providers—for example, by acquiring predeveloped system components, technological infrastructure (e.g., cloud computing abilities), or consultancy support: "I always like to understand how [potential providers] build their models, what they can do, who owns them or the training data sets, etc." (AI innovation manager).

Moreover, due to ethical, legal, security and performance issues, Siemens has to thoroughly check whether its collected use cases violate governance policy. Each use case project must state why it sufficiently aligns with all aspects of Siemens's governance framework. Doing this was especially challenging for the chatbot-based SiemensGPT platform. Given the system's large scale, key questions arose regarding global cybersecurity, cross-location human accountability, data sensitivity, etc.

Risk 2: Algorithmic Training and Data Quality Issues

Good task data is the backbone of every AI project. Without it, machine learning algorithms can be neither trained nor updated over time, meaning that data is a truly strategic resource.[26] Unfortunately, high-quality data from both internal and external sources is often scarce and not easily accessible to firms. In an industrial setting, firms have to install many types of sensors to access technical product or process data in very different locations. Moreover, when collecting customer or supplier data there may be sensitivity, privacy or legal restrictions. In addition, data is not always ready to use and often requires manual cleaning. For instance, datasets can be outdated, inconsistent, incomplete, unstructured or reflect biases regarding gender, race, age or educational background: "Building a [good] data foundation … is one of the biggest topics at the moment … for all these data analysts and our colleagues in various technology departments" (AI transformation strategist).

Another challenge Siemens faced resulted from the sheer size of the company. Even though one department may have access to high-quality data, others with similar needs often remain unaware of this data, resulting in isolated data silos. For many impactful use cases, the company has to search extensively to find and aggregate relevant training data from independent business functions: "For most of our use cases, we need [to collect] very specific data sets, like [market] data, product data, sales data, etc. … This can be very tough" (Head of digitization for a business unit).

Risk 3: Task-Specific System Complexities

The third risk category arises from the direct task environments for which an AI system is implemented. For instance, embedding sophisticated AI applications within other task-related technological infrastructures can be very complex. AI systems typically have many interfaces with other IT infrastructures, such as hardware, cloud servers, enterprise software or industry robots. In turn, inappropriate system integrations or task deployment could lead to task failures. Thus, Siemens's Industrial Copilot must be able to assist in coordinating multiple mutually interacting shop floor machines simultaneously: "The [technological task] landscapes are so complex. There is not just one integration with one other system. There are so many different ones, even at different locations" (AI transformation strategist).

If task conditions fundamentally change over time, AI systems already in use typically require appropriate algorithmic retraining. Outdated or inconsistent decision rules can significantly degrade system performance. For example, Siemens retrained an AI-based object classification system because of changes in lighting conditions within production facilities. Due to new types of light bulbs, the intelligent camera system had to partially relearn the task under different lighting conditions: "You need an ongoing process to adapt your machine learning model [when tasks change]. You need this way more than in a classic software system" (AI portfolio manager).

26 Hartmann, P. and Henkel, J. "The Rise of Corporate Science in AI: Data as a Strategic Resource," *Academy of Management Discoveries* (6:3), September 2020, pp. 359-381.

For many tasks, AI systems must be sufficiently understood and validated by technology experts before they can be exploited. In other words, firms must often guard against AI being perceived as a black box. This is particularly relevant for business-critical contexts, where single task failures can result in substantial negative consequences for the firm.[27] An example would be an ineffective or malfunctioning product design suggestion from the Virtual Prototyping tool that leads to erroneous machine product concepts. High change costs, investment losses or reputational damages could occur when the machine is put into production. Unfortunately, machine learning algorithms are probabilistic and learn massive amounts of complicated decision rules. Hence, even experts can not always fully retrace its task behavior. For Siemens, reducing black box perceptions is a critical but complex job: "[System validations are] often something that gets underestimated and take longer than anticipated" (AI system owner).

Risk 4: System Stakeholder Mismanagement

The fourth key risk category has to do with the management of relevant human and organizational system stakeholders. For each implementation project, a firm needs to identify all relevant individuals and entities inside and outside of the organization, as well as their specific needs. Alignment with these individuals and other organizations can help prevent misunderstandings or conflicts early on in the system development and implementation process. When Siemens developed the Virtual Prototyping AI product, it had to identify and consider the individual requirements of engineers, sales teams, customers and even regulators. Meeting stakeholder requirements often meant embedding specific technical functionalities, designing user-friendly interfaces, establishing reliable database connections, or ensuring clear system responsibilities: "You can generate the fanciest tool, but [stakeholders] often want completely different things that you have not on your mind. You have to gather feedback" (Data scientist/developer).

Furthermore, the Siemens case underlines the need to continuously identify and prevent potential governance violations over an AI system's lifetime. In practice, project contexts can significantly vary over time, potentially conflicting with enforced policies. Erroneous system updates or dynamic application contexts can create biased outputs, lead to ineffective system usage or create technology-rejection attitudes among stakeholders, strongly undermining scaling goals. Also, unforeseen cybersecurity risks may emerge when external system interfaces adapt, exposing the firm to hacking attacks and jeopardizing sensitive data.

Siemens also recognized the need to build AI systems that are understandable not only by experts but also by non-experts to ensure human trust and effective system usage. When stakeholders intensively interact with an AI agent, they usually need to understand both its decision-making behavior and its proper task application.[28] Such an understanding is particularly crucial if humans are personally affected by an algorithm; otherwise, they could develop an aversion to AI. When Siemens developed the Industrial Copilot, it was essential to make the system comprehensible to both employees and customers: "When you take [system stakeholders] along the journey, it helps to remove every fear, because then they understand they are with you. They can trust what you are doing [with AI]" (Chief data scientist).

Risk 5: Provider and System Dependencies

The fifth and final risk category concerns provider and system dependencies. Like most firms, Siemens collaborates with external technology providers to access infrastructure or knowledge competencies. For SiemensGPT, the company acquired pretrained large language models as well as cloud computing services

27 The research field of explainable AI explicitly deals with mitigating algorithmic black box perceptions. For an introduction to explainable AI, see Doshi-Velez, F. and Kim, B. "Towards a Rigorous Science of Interpretable Machine Learning," *arXiv*, March 2017, available at https://arxiv.org/abs/1702.08608.

28 For more details about the understandability needs of different human and organizational system stakeholders, see Meske, C., Bunde, E., Schneider, J. and Gersch, M. "Explainable Artificial Intelligence: Objectives, Stakeholders, and Future Research Opportunities," *Information Systems Management* (39:1), Spring 2022, pp. 53-63.

Figure 1: Five Critical AI Management Risks and Corresponding Mitigation Practices

The Five Critical AI Management Risks				
Risk 1: Missing or Falsely Evaluating Potential AI Opportunities	**Risk 2: Algorithmic Training and Data Quality Issues**	**Risk 3: Task-Specific System Complexities**	**Risk 4: System Stakeholder Mismanagement**	**Risk 5: Provider and System Dependencies**
How to Mitigate the Risk	**How to Mitigate the Risk**	**How to Mitigate the Risk**	**How to Mitigate the Risk**	**How to Mitigate the Risk**
1.1 Connect to your global AI ecosystem to stay abreast of technology trends	2.1 Implement internal data-sharing principles to create a central and current data warehouse	3.1 Observe virtual and physical system interdependencies to prevent task changes when possible	4.1 Integrate relevant stakeholders in the system development process step-by-step	5.1 Focus on multiple providers with comparable offerings
1.2 Establish a hub-and-spoke approach to search for promising AI use cases	2.2 Own critical data sources and protect them against potential imitators	3.2 Automatically retrain systems when conditions change	4.2 Constantly control and update company-wide policies	5.2 Steadily enhance your AI expertise through internal experimentation and external collaboration
1.3 Thoroughly discuss use case contexts, make-or-buy options and adoption constraints	2.3 Modularize system components and store them in an internal "model zoo" to future-proof the architecture	3.3 Classify tasks by their associated failure consequences to balance system interpretability and explainability	4.3 Educate stakeholders through workshops, technology promoters and system explanations	5.3 Let human experts learn from algorithmic task knowledge

with AI-based business analyst features[29] from vendors worldwide. The more critical an external firm becomes for an AI project, the greater the threat of potential lock-in effects—i.e., extraordinarily high costs for switching between providers. Lock-in effects generally make firms vulnerable to opportunistic provider behavior and limit their flexibility to take advantage of future technological changes, such as novel system releases from other firms. In the past, Siemens had faced lock-in threats when it fine-tuned sensitive data on large language models from only one provider: "When it comes to dependencies, we always need to be ready when a new provider, for instance, releases an AI functionality that the others do not have" (Principal AI technology expert)

In addition to provider dependencies, internal system dependencies might also limit a firm's flexibility to take advantage of future technological changes. In particular, Siemens is aware of so-called knowledge-loss risks when switching machine learning systems. Good algorithmic training data reflects detailed information about task contexts, process structures or factors influencing performance, potentially including knowledge that even human experts are unaware of. Switching could cause task knowledge to be lost if the training data of an existing and a former system are significantly different. Siemens identified knowledge-loss risks when it evaluated whether to replace an internally developed train control system with another one. Only the former included knowledge about train tracks that are no longer in use. Such knowledge could be advantageous for future infrastructure projects: "When you look at a machine learning system, it is not only the algorithm that is worth the money. It is the combination of the ingenuity in the algorithm plus the data set, and you have to combine the two" (AI transformation strategist).

Recommendations for Mitigating AI Management Risks

Having described the five critical AI management risk categories, we now provide three recommendations for handling each risk effectively. These recommendations, which are derived from Siemens's risk mitigation practices and summarized in Figure 1, unify both general technology and AI-specific management approaches. When implemented, they can

29 Many cloud and standard software providers offer additional AI features like business analytics or project management tools.

substantially help practitioners to better scale AI systems for their own business contexts.

1. Recommendations for Mitigating AI Opportunity Identification and Evaluation Risks

1.1. Connect to Your Global AI Ecosystem to Stay Abreast of Technology Trends

As a first step to not missing potential AI opportunities, Siemens identifies and reassesses its global AI ecosystem, including all relevant use case information sources. This approach underscores the need for a systematic approach because AI is much more diverse and dynamically advancing than many other technologies. Exponentially growing investments in AI from all industries, research institutions and governments underline a technological megatrend.[30] As a consequence, Siemens systematically divides its AI ecosystem into three separate layers to effectively perceive and interpret relevant technology signals.

The first layer comprises in-house departments such as R&D or IT functions, which can often provide signals about novel AI opportunities and trends. However, since it is not only employees with technology backgrounds that have valuable ideas, firms should use their entire company network to identify interdisciplinary pioneers in every department. To address this issue, Siemens established an "AI lab," an experimental setting in which employees from all its international divisions can develop and exchange ideas. The best concepts can then be presented to potential customers or at conferences: "[AI labs] are connecting data enthusiasts [and] different departments. They make sure that we understand what is happening in the … technological space" (Head of digitization for a business unit).

The second layer is the industry environment in which a firm operates. Siemens emphasizes the need to understand the business activities of key customers, suppliers and partners in order to develop innovative solutions for or with these stakeholders. It also assesses the initiatives of competitors to ensure it does not get left behind in strategic AI areas. In addition, Siemens seeks out opportunities to collaborate with all types of industry startups: "We have a strong startup network. ... We [want to see] which of those startups might be suitable to codevelop [with us]" (Head of digitization for a business unit).

The third ecosystem layer is outside the industry segments in which Siemens actively operates. This layer includes firms from non-engineering industries, consulting companies and research institutions. In particular, Siemens seeks out other technology leaders with innovative ideas that can be adapted and transferred to an industrial AI context: "For example, banks. ... Due to the digitalization needs in their industry, they are sometimes a bit more ahead. We [have already] had some really interesting exchanges" (AI transformation strategist).

1.2. Establish A Hub-and-Spoke Approach to Search for Promising AI Use Cases

Because of its large, global AI ecosystem, Siemens uses a hub-and-spoke model to scout for promising AI use cases. This model comprises a centralized department that generally takes care of technology ecosystems, together with decentralized teams within different business functions that build on their individual connections. Thus, although it has a large corporate technology scouting division, Siemens also encourages teams and individuals from diverse departments to connect with relevant ecosystem sources. Though this hub-and-spoke model focuses on all kinds of technologies, as the AI ecosystem constantly grows, more and more employees specialize in machine learning topics: "[AI experts] act as … a funnel. They take the time to look in the market, figure out what is happening, what is coming and if it could be useful for us" (Head of AI for procurement).

The hub-and-spoke model is most effective when it combines technology and task perspectives. The technology lens finds use cases for a specific type of AI application—natural language processing, computer vision, business analytics, etc. In contrast, a task lens assesses how predefined business contexts can be generally improved with machine learning. Combining both perspectives unifies general technology and specific business problem knowledge. A notable example is the Virtual Prototyping tool. To identify this industrial use case, the company

30 For details about the outstanding relevance of AI compared to other technologies, see Yee, L., Chui, M. and Roberts, R. *McKinsey Technology Trends Outlook 2025*, McKinsey & Company, July 22, 2025, available at https://www.mckinsey.com/capabilities/mckinsey-digital/our-insights/the-top-trends-in-tech.

needed both technology knowledge from AI experts as well as task expertise from engineers and developers: “Corporate technology scouts do the AI-specific front-end research … but they do not [necessarily] have an understanding of how we are working [in the functional departments]. You must directly connect both sides” (Head of AI for R&D).

Remarkably, a firm can even use AI agents to identify important information from its ecosystem. For instance, machine learning systems can automatically screen massive amounts of online news or press releases to derive innovative use case ideas.

1.3. Thoroughly Discuss Use Case Contexts, Make-or-Buy Options and Adoption Constraints

To determine which identified use cases qualify for a pilot phase, Siemens evaluates both business- and technology-specific factors. For example, prototyping with digital twins reduces internal resource needs but can also create innovative industrial customer offerings. The company also considers financial and productivity metrics, as well as the transferability of use cases to other similar task contexts. Siemens also assesses the technological feasibility of underlying machine learning models, which involves evaluating data collection needs, infrastructure requirements and realistic algorithmic accuracy levels. These metrics directly determine a system’s overall task performance: “The best way of doing [algorithmic evaluations] is to have a test scenario where you can have benchmarks. … Maybe it is a reduction of complexity, but I always look at benchmarks … to decide” (Head of technology for a business unit).

Interestingly, given the typically high investment costs for global AI implementation projects, Siemens emphasizes that not every use case requires a machine learning solution. Many proposed use cases can be addressed effectively with traditional information systems or simply by adjusting organizational processes. A notable example was the discussion about introducing an automated project-monitoring system to control software development progress. Instead of implementing an AI tool, Siemens changed task responsibility allocations and made project status meetings mandatory, leading to a measurable productivity increase: “I can tell you [more than] 50% are not AI use cases, which is fine. There are other technologies or ways” (Head of AI for a business unit).

In addition to assessing use case contexts, firms must decide whether to make or buy an AI solution. For example, critical training data, computing infrastructure or specific process knowledge might be externally acquired to reduce AI adoption times. Deciding whether to make or buy requires firms to be aware of their own technology competencies and those of available providers. Outsourcing also means anticipating potential supplier dependency threats. When outsourcing is seen as the best option, Siemens seeks to evaluate the competencies of potential providers without committing to contractual obligations too early—for example, through a proof of concept: “Our approach is to contact [providers] to get a trial license or something like that, to learn about focal components or services, to get familiar with them” (Data scientist/developer).

A final essential aspect of evaluating AI use cases is to consider adoption constraints, which are often related to the firm’s AI governance policies. If there will be significant problems regarding algorithmic accountability, cybersecurity, understandability, safety or privacy needs, projects must be rethought or withdrawn. For instance, AI-specific regulations like the European Union’s AI Act[31] might impose explicit technology adoption constraints. In fact, for some global projects, Siemens was unable to partner with certain providers because they did not fully comply with the European act: “The machine learning provider does not meet the [legal] requirements we want to have. We need to have a service in Europe ... but still, data can be processed in the U.S. … So, unfortunately, we cannot use these kinds of AI providers” (AI transformation strategist).

2. Recommendations for Mitigating Algorithmic and Data Quality Risks

2.1. Implement Internal Data-Sharing Principles to Create a Central and Current Data Warehouse

31 For more information about the European AI Act, see *EU AI Act: First Regulation on Artificial Intelligence*, European Parliament, June 8, 2023, available at https://www.europarl.europa.eu/topics/en/article/20230601STO93804/eu-ai-act-first-regulation-on-artificial-intelligence.

To improve the quality of and access to internal training data, Siemens has established very effective data-sharing principles to ensure that individual data scientists from different departments proactively share their datasets and keep them current over time. The data scientists are incentivized, for example, through bonus payments or new technology equipment, to steadily collect and share domain-specific product or process data while observing the firm's governance guidelines:[32] "One of the principles that we are pushing for is what we call data-sharing principles. Share unless there are legal regulations or contractual limitations as to why you cannot share data" (Chief data scientist)

The ultimate goal for Siemens is a fully centralized data warehouse that consolidates information from all business functions worldwide. With such a central data pool, AI experts can quickly identify and extract relevant training data to create or update their own algorithms. This was one of the reasons for the company-wide rollout of a common data cloud, facilitating the storage of data at scale and sharing it in a governed way within the organization. In the future, SiemensGPT could also be used to search for specific AI datasets or applications within the global company network and, if necessary, approach those responsible for the datasets: "It is a platform way where things are being listed and show what databases, what information is available inside the company. … [We have] a search engine to help people find suitable databases and the corresponding contact persons" (Data scientist/developer).

2.2. Own Critical Data Sources and Protect Them against Potential Imitators

Because data has become a strategic resource, Siemens aims to own critical data sources and protect them internally. As a first step, it classifies the business relevance of training data. For instance, Siemens expanded its sensor landscape to better measure the physical properties of machinery during R&D tests, enabling more effective algorithmic training data for the Virtual Prototyping system. The most cost-efficient data-acquisition approach is selected for less relevant data, which may involve collecting data internally, purchasing data or accessing it from open-source platforms: "It is not just about availability; it is also about quality. … How critical are certain datasets for your algorithms? The [critical datasets] are the ones you must own. No way around that" (AI transformation strategist).

Critical data that is owned must be properly protected to prevent imitators from building similar systems or training even better ones, resulting in the loss of technological advantages. Good protection not only consists of intellectual property measures but also involves state-of-the-art cybersecurity standards. Because many industrial AI offerings are based on high-quality product data, Siemens installed several alarm mechanisms to detect suspicious anomalies within internal data streams. The foundation of this approach is the zero-trust concept, which assumes that fraud and human errors can occur anywhere in an IT landscape.

However, critical data is not always available in-house, meaning that Siemens must frequently access external sources. The company must make sure that external data quality, quantity and compatibility will be sufficient and consistent over time. Therefore, Siemens seeks to secure multiple independent external data sources whenever possible. Combining data from different external sources can also result in higher data quality, thus providing resource advantages. Siemens also pools data with other firms. For example, it partners with suppliers to enable just-in-time deliveries, cooperates with customers to personalize services and bundles research data with startups and universities.

2.3. Modularize System Components and Store Them in an Internal "Model Zoo" to Future-Proof the Architecture

Siemens actively promotes modularity in the algorithmic components of its AI systems to streamline future training processes. Indeed, when an algorithm consists of multiple model-building blocks, it can simply replace or retrain certain parts instead of the entire system. Moreover, modularity enables colleagues around the globe to configure new types of applications or enhance existing ones with additional training data. To achieve modularity, Siemens uses a so-called "model zoo"—a virtual model storage in which experts can share individual components

32 For an insightful article on data-driven firm culture, see Staudt, P. and Hoffman, R. "How a Utility Company Established a Corporate Data Culture for Data-Driven Decision Making," *MIS Quarterly Executive* (23:1), March 2024, pp. 19-35.

of their AI systems and the underlying training data. Building on these pretrained model blocks can significantly reduce development times and data complexities. Experts searching for specific algorithmic functionalities can directly access suitable model components that have already been stored: "When you start to build an application, you need to directly think, what are the different parts in my architecture, and build the architecture in a way that you can swap those parts" (Head of AI for R&D).

3. Recommendations for Mitigating Task-Specific Complexity Risks

3.1. Observe Virtual and Physical System Interdependencies to Prevent Task Changes When Possible

One essential aspect of mitigating task-specific AI risks is proactively identifying and evaluating potential system interdependencies within the applied adoption context. First, such interdependencies can result from virtual interfaces to other digital systems that deliver input, use output, work in parallel or provide complementary infrastructure services. Second, there are physical interfaces, such as when humans engage in coworking with robots. In this case, aligning algorithmic and natural body movements is necessary. Alternatively, as mentioned earlier, lighting conditions, radiation or other physical factors can influence algorithmic task performance. As a result, Siemens has started systematically screening task environments to anticipate and prevent potentially disruptive changes, with the aim of ensuring a functioning status quo whenever possible: "You need to look into the complete [task] environment. The model itself is just one building block for a [successful task performance]" (Head of digitization for a business unit).

3.2. Automatically Retrain Systems When Task Conditions Change

Even with the most significant prevention efforts, task conditions can still change substantially. This typically means that an existing AI system must be retrained to adapt its previously self-learned decision rules to changed task conditions. Ideally, there should be an automated pipeline feeding an algorithm with the latest task data. Siemens calls this the "closed-loop approach." An illustrative example is provided by Siemens's financial sales department, which applies predictive AI to estimate customer payment defaults. For instance, the AI system is able to update itself in response to unseen contractual obligations. "In our customer credit decision process, we have a predictive machine learning model. … The whole project was already set up to have model monitoring in place, and dedicated retraining possibilities that … make it not that much effort … when new customer data is available" (AI transformation strategist).

3.3. Classify Tasks by Their Associated Failure Consequences to Balance System Interpretability and Explainability

Another issue related to task-specific complexity risks is that humans perceive many complex and probabilistic machine learning-based AI systems as black boxes. However, experts must sufficiently understand and validate many of these systems before they can be implemented. Hence, as a starting point, Siemens determines the required level of expert understandability based on the potential task impact of system failures. AI adoption tasks are classified into higher- and lower-risk categories. For instance, flawed supply chain predictions can lead to massive production shutdowns, signaling a higher risk level. The same holds for industrial customer products because malfunctioning systems can result in reputational damage or even legal disputes. In contrast, creating product advertisement ideas with generative AI tools is usually much less critical. In sum, the higher the risk category for a task, the higher the understandability requirements to validate an algorithm. "If you apply self-learning AI models, you always ask yourself, what is the real risk of a wrong answer [by the system]?" (AI innovation manager)

To achieve the required level of expert understandability, firms have to carefully balance two approaches. The first approach focuses on realizing system interpretability—i.e., creating an inherently understandable AI model that is truly transparent for technology experts. System interpretability is particularly relevant for critical tasks because it allows experts to largely retrace an algorithm's self-learned decision rules. For the Virtual Prototyping use case, Siemens focused on achieving interpretability because erroneous design suggestions could have negatively affected long-term R&D projects. For that reason, system developers steered algorithmic training processes and marked the most relevant task variables. Alternative interpretability measures might involve the manual analysis of datasets or statistically simpler training methods.

Given the cognitive limitations of humans, system interpretability could conflict with an algorithm's self-learning power.[33] Hence, the second approach is to implement system explainability, which provides additional explanations in the form of texts, audio or images by using complementary generative AI tools. These artificially created explanations can serve as an impactful control mechanism because they can inform experts about a system's decision-making tendencies. Nevertheless, while explainability only approximates the behavior of an algorithm with another one, it does not really solve the black box issue, making it more suitable for less critical task contexts.

4. Recommendations for Mitigating System Stakeholder Mismanagement Risks

4.1. Integrate Relevant Stakeholders in the System Development Process Step-by-Step

Because AI systems must typically meet diverse human and organizational stakeholder needs, Siemens sequentially integrates important internal and external stakeholders in the system development process. Such stepwise codevelopment helps developers and stakeholders to better understand what the final system should look like and fosters an interdisciplinary AI competency within the entire firm. To facilitate this practice, Siemens creates a question catalog for use case projects. Based on the answers given, relevant stakeholders are integrated into the project at the appropriate stage. For example, cybersecurity specialists get involved when an AI application has access to external systems from suppliers, customers or ecosystem partners. When developing the Industrial Copilot, Siemens sequentially integrated experts from its production facilities who performed a "customer-zero" role, providing potential recommendations from a market perspective: "We infused the [existing compliance and development] process with trigger questions here and there. So, depending on what you answer, the correct team is involved to help and guide you … before launching a system [internally] or to a customer" (Principal AI technology expert).

4.2. Constantly Control and Update Company-Wide Policies

Because compliance violations could lead to undesirable consequences, Siemens established human accountability structures to continuously control a system's alignment with enforced policies. Specifically, it allocates clear responsibilities to both system and task owners. If violations are found, the stepwise integration of affected stakeholders might start again until all threats are eliminated: "We [have] a continuous compliance process with clear roles and responsibilities. … We need to ensure that our developments and deployments [always] follow the proposed governance framework" (Principal AI technology expert).

Siemens regularly reassesses and updates its governance framework to take account of changing or emerging risks. An example is governance changes in response to new cybersecurity threats driven by deepfakes. The company has experienced an increasing number of deepfake-based phishing attempts, including deceptive voice messages or manipulated images, to steal company-sensitive information.

4.3. Educate Stakeholders through Workshops, Technology Promoters and System Explanations

Firms should provide education for non-expert AI system stakeholders so that they can better understand and use AI systems. Siemens regularly organizes workshops about the benefits of using AI technologies and how to use them. At these workshops, participants can learn about specific AI systems in a playful environment featuring exercises and experiments. Those who successfully complete the workshop receive AI certificates, underlining their newly gained skills: "The important thing is that [non-experts] get some hands-on experience within a playground setting. … I think trust comes with a lot of interaction and understanding of what to expect.

33 The understandability-performance trade-off is further described in Rai, A. "Explainable AI: From Black Box to Glass Box," *Journal of the Academy of Marketing Science* (48:1), December 2019, pp. 137-141.

We are very strongly promoting such AI literacy courses" (Principal AI technology expert).

Siemens also empowers more and more so-called AI ambassadors or promoters to enthusiastically foster the exploitation of its implemented AI systems. Ambassadors and promoters are employees who voluntarily advocate the use of AI tools, encouraging their colleagues to use them for daily tasks. In other words, AI ambassadors and promoters function as personal contacts for those who have open questions or concerns. For example, ambassadors recently convinced colleagues to use the central SiemensGPT platform instead of individual AI applications: "You have people who are interested in technologies and try to work with new [AI systems]. ... If you would like to implement a new technology, you should always focus on these kinds of users first because they can be [AI] ambassadors" (Head of technology for a business unit).

Another way to educate non-expert stakeholders is to create direct system explanations.[34] In contrast to expert explanations that help validate a system's correctness, explanations for non-experts should describe a system's general purpose and behavior to foster trust. For instance, explanations could describe why an algorithm came up with specific outputs or why it is unbiased. Customers could ask the Industrial Copilot why certain code may or may not be successful. Specifically, the system could analyze code snippets and verbally explain potential error sources.

Explanations do not necessarily have to be in a text format. Siemens's Virtual Prototyping tool provides diagrams to help engineers understand simulation results. However, explanations must always be expressed in an intuitively comprehensible manner considering the social needs of stakeholders—for example, by using non-technical language and context-specific examples:[35] "You have to increase trust from a social point of view ... making [system stakeholders] believe that [the AI system] is not confabulating, that it sticks to the topic" (Data scientist/developer).

5. Recommendations for Mitigating Provider and System Dependency Risks

5.1. Focus on Multiple Providers with Comparable Offerings

To avoid critical provider dependencies, Siemens has adopted a dual-sourcing approach for impactful projects, partnering, wherever possible, with at least two comparable providers. For instance, it has integrated multiple large language models from different providers into the SiemensGPT platform, allowing users to choose among models. This approach also supports replacing or retraining specific models while ensuring uninterrupted operation. In addition, Siemens emphasizes the importance of standardized application programming interfaces (APIs) when connecting external and internal technology systems. These API standards enable uniform communication between different types of digital systems, facilitating switching efforts: "[We] always [...] have an adaptation layer, like an API, that is similar or comparable so that the programming interface does not need to be adjusted all the time as soon as a new model is released or [providers] are changed" (Head of technology for a business unit).

5.2. Steadily Enhance Your AI Expertise through Internal Experimentation and External Collaboration

The more AI expertise a firm develops internally, the less it must depend on external support. Siemens adopts a twofold approach to steadily advancing its technological knowledge and competencies. First, the firm fosters experimentation because it provides valuable learning for the future. Second, when collaborating with competent AI service providers, it tries to absorb new knowledge. A good example is the company's recruiting department, which needed specific support from external consultants to ensure the way it processed applicant data complied with European data protection regulations. However, since Siemens's legal experts has continuously expanded their AI knowledge, the recruiting department no longer needs external consultants: "Our legal colleagues [...] they are getting our consultants" (AI transformation strategist).

34 Often, expert explanations can be transferred into non-expert explanations.

35 For more insights on human-centered system explanations, see Miller, T. "Explanation in Artificial Intelligence: Insights from the Social Sciences," *Artificial Intelligence* (267), February 2019, pp. 1-38

5.3. Let Human Experts Learn from Algorithmic Task Knowledge

To mitigate knowledge-loss risks, Siemens recognizes that task experts can learn from the decision rules stored within machine learning algorithms.[36] In fact, experts can extract important information about task processes, influence factors and error sources. To do so, they use additional generative AI tools to ask task-specific questions to the focal algorithm. Individuals then reflect on the findings and exchange the gained insights. For example, Siemens's sales managers learn about the individual preferences of customers when an algorithm identifies patterns in their behavior. Nonetheless, it is always important to comply with intellectual property constraints, particularly for external collaborations: "You should write down [extracted] rules, not very specifically for one model, but rather in a neutral way so that you can reuse a catalog of rules" (Head of AI for R&D).

Another interesting observation from the Siemens case is that AI systems can mutually learn from each other. Though this concept is still in its infancy, it builds on algorithmic imitation—i.e., an AI system acquires knowledge from another system by statistically approximating its decision rules through imitating data input-output pairs. For example, regarding a use case in AI robotics, pretrained shop floor machines could further improve their performance by replicating actions that other robots have already learned, such as grasping material components or smoothly moving along assembly lines.

Concluding Comments

Today's machine learning-based AI systems provide a host of business opportunities. Nevertheless, due to the complex technological characteristics of machine learning, there are many critical technology management risks that firms must effectively mitigate to scale AI globally and gain measurable business impacts. To mitigate these risks, practitioners must combine general and AI-specific technology management practices.

To provide practical advice on managing and strategically scaling today's AI systems, we conducted an in-depth case study of Siemens AG, which has gained strong industrial AI competencies over time. During its successful AI journey, Siemens encountered five different but equally critical technology management risk categories that had to be mitigated:

1. Missing or falsely evaluated potential use case opportunities
2. Algorithmic training and data issues
3. Task-specific system complexities
4. Stakeholder mismanagement
5. Threats from provider and system dependencies.

Based on Siemens's experiences and current best practices, we provide three recommendations for mitigating each risk category. We encourage managers from both larger and smaller firms to adopt our recommendations to better scale generative and predictive AI systems. However, it is important to emphasize that the field of AI is advancing dynamically. Given the rapidly changing technology, firms should constantly reassess their own AI landscapes and applied technology management practices to ensure that they can strategically scale machine learning-based AI systems.

Appendix: Research Method

Our research for this article comprised a single in-depth case study of Siemens AG for the purpose of intensively exploring two research questions.[37] This company was selected because of its strategically impactful AI projects and strong industrial AI competencies. The primary data source was 21 interviews with various Siemens technology experts and strategists (see table), but we also analyzed several of the firm's documents and correspondence regarding implementation processes, technological

36 For a well-argued outlook on how machine learning can reshape organizational learning in the future, see Balasubramanian, N., Ye, Y. and Xu, M. "Substituting Human Decision-Making with Machine Learning: Implications for Organizational Learning," *Academy of Management Review* (47:3), July 2022, pp. 448-465.

37 Excellent instructions on conducting case study research are given in: 1) Yin, R. K. *Case Study Research and Applications: Design and Methods* (6th edition), SAGE Publications, 2018; and 2) Lee, A. S. "A Scientific Methodology for MIS Case Studies," *MIS Quarterly* (13:1), March 1989, pp. 33-50.

Interviewees' Roles within Siemens

No.	Role	Years of Relevant Experience
1	AI Innovation Manager	15 - 20
2	AI System Owner	5 - 10
3	Head of Technology for a Business Unit	15 - 20
4	Chief Data Scientist	15 - 20
5	AI Innovation Manager	> 20
6	Principal AI Technology Expert	10 - 15
7	Data Scientist/Developer	2 - 5
8	Head of Digitization for a Business Unit	10 - 15
9	AI Transformation Strategist	5 - 10
10	AI Transformation Strategist	10 - 15
11	AI Transformation Strategist	10 - 15
12	Head of Digitization for a Business Unit	5 – 10
13	AI Portfolio Manager	10 – 15
14	Data Scientist/Developer	2 – 5
15	Data Scientist/Developer	2 – 5
16	Head of AI for R&D	15 – 20
17	Principal AI Technology Expert	5 – 10
18	AI Transformation Strategist	5 – 10
19	Head of AI for Procurement	10 – 15
20	Head of AI for a Business Unit	5 – 10
21	AI Transformation Strategist	> 20

functionalities and management initiatives to better understand the case context.[38]

All interviewees were involved in globally scaled industrial AI projects, including both predictive and generative AI systems. We included suitable candidates from diverse corporate and business-level functions to obtain multiple organizational perspectives. The interviews were mostly conducted remotely during September and October 2024 and lasted approximately 50 minutes on average.

We followed a semi-structured interview guide that enabled us to be flexible in addressing upcoming thoughts and follow-up questions while following a systematic scheme.[39] This approach helped us to focus on the significant areas of expertise of each interviewee. We started with questions about the interviewee's background, ongoing AI endeavors at Siemens and substantial technology management risks. Next, we asked about practical risk mitigation approaches, leaving time for interviewees to provide additional thoughts and recap their most interesting arguments.[40] We stopped the interview process when high response

38 For additional information about data sources for case study research, see Benbasat, I., Goldstein, D. K. and Mead, M. "The Case Research Strategy in Studies of Information Systems," *MIS Quarterly* (11:3), September 1987, pp. 369-386.

39 To learn more about the purpose of interview guides, see Myers, M. D. and Newman, M. "The Qualitative Interview in IS Research: Examining the Craft," *Information and Organization* (17:1), January 2007, pp. 2-26.

40 More details can be found in Schultze, U. and Avital, M. "Designing Interviews to Generate Rich Data for Information Systems Research," *Information and Organization* (21:1), January 2011, pp. 1-16.

redundancy occurred because no more insights were revealed.[41]

After transcribing the interviews, we analyzed the data using open, axial and selective coding techniques from grounded theory research in multiple steps.[42] To avoid subjective biases, the authors iteratively discussed the findings of each coding round together.[43] This process allowed us to aggregate emerging codes to systematically detect the most critical AI management risks for Siemens, with respect to the specific technological characteristics of machine learning. Next, we evaluated and matched potential mitigation approaches for each identified risk category, triangulating our results with insights derived from various internal documents and practical system demonstrations. Finally, we shared our findings with Siemens to verify their validity and ensure real-world impact.

About the Authors

Thomas Hutzschenreuter

Professor Thomas Hutzschenreuter (th.sim@tum.de) holds the Chair of Strategic and International Management at the Technical University of Munich (TUM). His internationally awarded research focuses on strategizing, technology-enabled business and corporate strategies, organizational transformation, governance structures and firm growth. He has published in leading international journals, including *Strategic Management Journal*, *Journal of Management Studies*, *Journal of Management*, *Journal of International Business Studies*, *Leadership Quarterly* and *Academy of Management Perspectives*. Thomas is frequently invited to give keynotes and guest speeches and has received several Best (Executive) Teaching awards.

Tim Lämmermann

Tim Lämmermann (tim.laemmermann@tum.de) is a doctoral research associate at the Chair of Strategic and International Management at the Technical University of Munich (TUM). His research predominantly focuses on the intersection of strategic management and artificial intelligence, including AI-oriented business strategies, technology management practices and firm transformations. He regularly collaborates with executives and technology strategists from leading global firms and teaches AI strategy topics in university classrooms. In addition to his research and teaching, Tim is a senior AI strategist at a German management consultancy, where he systematically supports firms in their strategic AI transformations.

Alexander Sake

Alexander Sake (alexander.sake@siemens.com) is the head of finance at Siemens's Global Rail Infrastructure Project Execution and Engineering Hub. He also serves as an internal expert on the strategic use of AI in finance and commercial functions. With over 15 years of international experience across Germany, the U.S., Malaysia and India, he has co-led AI and digitalization initiatives, including transformation projects in the Asia-Pacific region. He holds an executive MBA from the Technical University of Munich, with a focus on strategy, innovation and sustainability.

Helmuth Ludwig

Helmuth Ludwig (hludwig@smu.edu) is a professor of practice in Strategy and Entrepreneurship, the executive director of the Hart Institute for Technology, Innovation and Entrepreneurship at Southern Methodist University in Dallas, Texas, and an executive fellow at IESE in Barcelona. Previously, he spent 30 years at Siemens in different executive roles, including as president of PLM Software and the CEO of Siemens' North America's Industry Business. Until 2019, he was the CIO of the Siemens Group. He also serves on the corporate board of Hitachi Ltd, chairs the board of the Humanetics Group and is a senior advisor with Bridgepoint, LLC (private equity).

41 See: 1) Corbin, J. M. and Strauss, A. L. *Basics of Qualitative Research: Techniques and Procedures for Developing Grounded Theory* (4th edition), SAGE Publications, 2015; and 2) Patton, M. Q. *Qualitative Research & Evaluation Methods: Integrating Theory and Practice* (4th edition), SAGE Publications, 2015.

42 See: 1) Gioia, D. A., Corley, K. G. and Hamilton, A. L. "Seeking Qualitative Rigor in Inductive Research: Notes on the Gioia Methodology," *Organizational Research Methods* (16:1), January 2013, pp. 15-31; and 2) Corbin, J. M. and Strauss, A. "Grounded Theory Research: Procedures, Canons, and Evaluative Criteria," Qualitative Sociology (13:1), March 1990, pp. 3-21.

43 Klein, H. K. and Myers, M. D. "A Set of Principles for Conducting and Evaluating Interpretive Field Studies in Information Systems," *MIS Quarterly* (23:1), March 1999, pp. 67-93.