Michael Lively-Scholz
Master's Thesis: CYSE 698
Advisor: Dr. Sultuk Karahan
8/2/26

Secure-by-Design AI Workflows: The AI Workflow Lifecycle Framework for Managing AI as an Organizational Capability

I. Abstract

Enterprises are rapidly integrating artificial intelligence into their operations and business processes. However, existing guidance for organizations primarily address artificial intelligence risk management or cybersecurity as opposed to addressing the entire organizational lifecycle requires to operationalize AI workflows. This thesis proposes the AI Workflow Lifecycle Framework (AWLF), a conceptual framework that outlines how an organization should manage AI workflows as organizational capabilities instead of software projects. This paper makes three primary contributions. It establishes a working academic definition of AI workflows, the AI Workflow Lifecycle Framework, and provides a secure-by-design operational philosophy for enterprise AI adoption. The framework positions "Business Need" as the organizational trigger for enterprise AI adoption. It emphasizes executive sponsorship, security, governance, and continuous lifecycle management. The AWLF provides organizations a lifecycle management solution that is intended to complement existing frameworks such as the National Institute of Standards and Technology's Cybersecurity Framework 2.0 (NIST CSF 2.0) and their AI Risk management Framework. The AWLF primarily outlines organizational capabilities that will improve the success of AI workflow adoption.

II. Introduction

Artificial Intelligence (AI) in recent years has evolved exponentially. Only a few years ago, it was an emerging, novel technology, but now it can be a foundational component of modern operations within an organization. Most enterprises are implementing AI into operations, and those that do not may soon be left behind. Across industry, AI is being implemented into software development, cybersecurity, healthcare, finance, manufacturing, and customer service to automate repetitive tasks, enhance informed decision making, and improve enterprise efficiency. Although AI, specifically generative AI, is a relatively new technology, it has allowed organizations to automate complex tasks that require prior knowledge. These jobs or tasks typically required an individual to have a college degree or significant industry experience in an area to have the necessary background knowledge required for the task. Now an individual can utilize AI to complete these tasks autonomously without much knowledge on the specific area.

As Large Language Models (LLM's) have become mainstream, companies that once offered a paid knowledge base for specific problems have felt the impact directly. Chegg, an online education company built to offer paid homework help and study guide solutions, experienced a stock decline of 25% after a February 2023 SEC filing referencing ChatGPT, and by almost 50% following a May 2023 filing in which the company said that rising ChatGPT adoption was hurting its new customer growth rate (Pietrzak, 2025). The next movement will likely involve companies that implement AI directly into workflows that support critical organizational processes. AI workflows are structured sequences of automated, repetitive jobs that utilize AI models to operate on data sources and enterprise applications combined with humans to make critical decisions in order to accomplish a specific objective that supports business strategy. This augments the work of an employee, it can serve as an intern, entry-level, or even in some cases mid-level employee, allowing the employee can focus on other higher-level tasks and act as a manager or decision maker. Workflows such as document processing pipelines, cybersecurity investigations, coding assistants or agents, and customer service agents have allowed organizations to reduce operational costs and improve productivity and efficiency. As technology continues to develop and solidify its niche inside enterprise, it is expected that AI will move from being utilized for small, occasional tasks to larger enterprise operational process adoption.

With the added value that AI workflows drive into organizations, they also bring forth a plethora of challenges. An organization cannot simply implement AI into their systems. It unleashes a host of cybersecurity and overall technical challenges to the organization. As opposed to traditional enterprise software, AI depends on probability models, evolving knowledge bases, third-party services, autonomous decision making, and dynamic, ambiguous prompts. This unique, inherent nature of AI models creates challenges for organizations, as they have to protect systems against poisoning of data the model uses to operate, prompt injection, supply chain vulnerabilities, and accidental leakage of sensitive or confidential information. These unique challenges require an organization choosing to implement AI within the enterprise environment to adapt and innovate traditional processes regarding the governance, security, and design process of their systems.

When an organization wants to utilize AI in their business and operational systems and processes to increase revenue, profits, or efficiency, they should not simply deploy a large language model (LLM) into their systems. Proper organizational AI implementation includes properly structured AI models that improve workflows

and are treated as enterprise software systems that require secure architecture, continuous governance, and built-in operational security from deployment. The successful adoption of AI within an enterprise requires organizational security ownership, involving executive leadership, along with strong engineering, governance, and lifecycle management. This, along with the use of a secure-by-design implementation approach, allows an enterprise to maximize the business value of AI while ensuring the organization experiences a reduction in risk to their operational and cybersecurity posture. Enterprise AI workflows should be managed as organizational capabilities rather than software projects. To operationalize this shift, this paper proposes the AI Workflow Lifecycle Framework (AWLF), a conceptual model presented in Section V and framed around enterprise organizations, though its principles may be utilized by organizations of any size or sector pursuing AI adoption. To support the framework, this paper presents a working definition of AI workflows and incorporates a secure-by-design philosophy for successful AI implementation. The proposed framework defines the necessary organizational capabilities that would contribute to successful AI workflow adoption for an organization rather than detailing specific security controls or technical implementations.

III. Background and Foundational Concepts

A. Artificial Intelligence

Artificial intelligence (AI) is a broad field of computer science focused on developing systems capable of performing tasks that typically require human intelligence, including learning, reasoning, problem-solving, perception, and decision-making (Russell & Norvig, 2021). The National Institute of Standards and Technology (NIST) defines AI as "a machine-based system that can, for a given set of human-defined objectives, make predictions, recommendations, or decisions influencing real or virtual environments" (NIST, 2023). Modern AI encompasses numerous disciplines, including machine learning, natural language processing (NLP), computer vision, and robotics, all of which contribute to the development of intelligent systems that will be able to support increasingly complex business operations. This same capability that makes AI valuable to an enterprise, the ability to make decisions and predictions on its own, is also what makes AI more vulnerable that traditional systems.

B. Generative Artificial Intelligence

Generative artificial intelligence (GenAI) is a subset of AI designed to generate new content rather than simply analyze or classify existing information. NIST AI 600-1 defines generative AI as a class of AI models that emulate the structure and characteristics of training data to produce synthetic outputs, including text, images, audio, video, and software code (NIST, 2024). The rapid advancement of generative AI has enabled organizations to automate knowledge-intensive tasks, improve productivity, and improve decision-making across a wide range of industries. There are many kinds of GenAI. However, Large Language Models are the most widely used in enterprise systems, which is why they are covered next.

C. Large Language Models

Large language models (LLMs) are transformer-based foundation models trained on massive collections of text to understand and generate human language (Vaswani et al., 2017). Rather than retrieving predefined responses, LLMs generate text by predicting the most probable sequence of tokens based on learned statistical relationships within their training data (Brown et al., 2020). These models serve as the foundation for many enterprise generative AI workflows or applications, powering capabilities such as conversational assistants, document summarization, code generation, and intelligent knowledge retrieval. Bommasani et al. (2021) describe these systems as foundation models because they can be adapted to a wide variety of downstream tasks, making them a foundational technology for modern enterprise AI applications. As organizations increasingly integrate LLMs into business processes, they have become the primary reasoning engine behind enterprise AI workflows. Because LLMs generate output probabilistically rather than deterministically, this reliance introduces risks that traditional software does not carry, which is the subject of the next section.

D. AI Risk Management

The widespread adoption of generative AI has introduced new technical, operational, and cybersecurity risks that extend beyond those associated with traditional software systems. Recognizing these challenges, the National Institute of Standards and Technology developed the Artificial Intelligence Risk Management Framework (AI RMF) to help organizations identify, assess, and manage AI-related risks throughout the AI lifecycle (NIST, 2023). The framework is organized around four core functions: Govern, Map, Measure, and Manage. This provides organizations with a structured approach for implementing trustworthy AI while promoting transparency, accountability, and continuous risk

management (NIST, 2023). The framework is further expanded in NIST AI 600-1, the Generative Artificial Intelligence Profile, which addresses risks specific to foundation models and large language models (NIST, 2024). These frameworks provide the governance foundation this thesis builds on but applying them at enterprise scale brings its own set of adoption challenges, covered next.

E. Enterprise AI Adoption

Organizations across both the public and private sectors are rapidly integrating AI into business operations to improve productivity, automate knowledge-intensive tasks, and accelerate organizational innovation. While many have successfully demonstrated AI capabilities through pilot projects and limited deployments, scaling these technologies across enterprise environments remains a significant challenge. Integrating AI into production systems requires more than deploying an individual model; it demands robust architecture, governance, cybersecurity, and ongoing operational support. As enterprise AI adoption continues to mature, organizations are increasingly shifting their focus from standalone AI applications to integrated AI workflows that securely connect models, enterprise data, business applications, and human decision-makers to automate end to end business processes (NIST, 2024; McKinsey & Company, 2025).

F. AI Workflows

Although the concept of AI workflows is increasingly referenced throughout industry, there is currently no universally accepted academic definition describing what constitutes an AI workflow. Consequently, this paper adopts the following working definition: An AI workflow is a structured AI-enabled system that has a structured role and can reason, ideate, or make contextual decisions within a specific set of parameters to contribute towards a larger goal; AI workflows should assist humans in completing larger tasks. By completing a task in an autonomous manner, they can take ownership of a specific subset of a larger organizational process. An AI workflow ultimately assumes responsibility for a repeatable task while collaborating with humans, software systems, or other AI components to achieve a defined objective.

Unlike standalone AI tools like ChatGPT or other standalone LLM's, AI workflows use AI to automate repeatable organizational processes in an effort to support larger business needs or efforts. AI workflows often combine AI models, like LLM's or other forms of GenAI, with data sources, business applications, automation platforms, and human oversight to conduct knowledge-intensive tasks within business processes.

Enterprise AI workflows are different from traditional software automation as it incorporates probabilistic decision-making models capable of adapting to varying inputs instead of following exclusively deterministic business logic. Organizations implementing AI workflows must address unique considerations related to governance, lifecycle management, and cybersecurity that extend beyond traditional software development practices.

The increasing adoption of enterprise AI workflows shows a drastic change in how organizations utilize AI. This transition introduces new technical and organizational responsibilities that motivate the AI Workflow Lifecycle Framework proposed in this thesis.

IV. Business Need

AI adoption should begin with a validated business need, not technological enthusiasm or as a measure to stay relevant with business or enterprise trends. Within the AI Workflow Lifecycle Framework, AI adoption begins with a validated business need rather than technological enthusiasm or competitive pressure. Artificial intelligence has rapidly evolved from an emerging technology into a strategic organizational capability. Organizations across both the public and private sectors are increasingly implementing AI to improve operational efficiency, reduce costs, automate repetitive tasks, enhance customer experiences, and support data-driven decision-making. Although these capabilities offer significant organizational value, AI should not be adopted solely because it aligns with industry trends or pressures. Organizations should first establish a validated business need demonstrating how AI will advance organizational strategy, create measurable value, and justify the long-term investment required for successful adoption. AI workflows should be implemented with clearly defined objectives, organizational ownership, and measurable performance expectations, much like any other critical organizational capability.

The purpose of the Business Need stage is not simply to determine whether artificial intelligence can create value, but whether artificial intelligence is the appropriate solution to a specific organizational problem. A validated business need should identify a measurable operational challenge and demonstrate how an AI workflow can improve organizational performance while supporting broader business strategy. This distinction ensures that AI adoption remains driven by organizational objectives rather than technological enthusiasm.

Organizations increasingly pursue AI because it offers substantial opportunities to improve productivity and operational efficiency. Coyne et al. (2019) estimate that approximately 50% of workplace activities are technically automatable using existing technologies, although relatively few occupations can be completely automated. Rather than replacing workers entirely, AI is expected to augment employees by automating repetitive tasks while enabling personnel to focus on higher-value activities requiring creativity, critical thinking, and strategic decision-making. Beyond individual productivity improvements, AI is projected to generate between $3.5 trillion and $5.8 trillion in annual economic value across major industries, representing nearly 40% of the value created through advanced analytics techniques (Manyika & Sneader, 2018). This illustrates why organizations increasingly view AI as a strategic investment capable of producing significant operational and financial benefits.

The business value of AI is particularly evident within software engineering. Industry studies demonstrate that generative AI coding assistants significantly reduce the time required for routine development activities. Deniz et al. (2023) found that developers completed documentation tasks in approximately half the traditional time, while substantially reducing the effort required for code generation and refactoring. Similarly, Peng et al. (2023) reported that developers using GitHub Copilot completed programming tasks 55.8% faster than developers working without AI assistance. These productivity gains illustrate how AI workflows can reduce operational costs while enabling employees to focus on higher-value engineering activities that require human judgment and creativity.

While these productivity improvements provide a compelling rationale for AI adoption, anticipated value alone does not constitute sufficient justification for implementation. Organizations should establish measurable objectives and expected returns on investment (ROI) before committing significant financial and organizational resources to AI initiatives. A validated business need should therefore identify quantitative performance indicators demonstrating how AI implementation is expected to improve operational efficiency, reduce costs, increase revenue, improve service quality, or otherwise advance organizational strategy. AI initiatives should be evaluated using measurable business outcomes rather than assumptions regarding the capabilities of emerging technology.

A validated business need must also align with the enterprise's broader strategic objectives. Rather than funding isolated AI initiatives that duplicate capabilities or compete for organizational resources, organizations should develop a comprehensive strategy that coordinates AI adoption across business functions. AI workflows should be designed with clearly defined organizational roles, responsibilities, and success criteria rather than deployed as standalone technology projects. Aligning AI initiatives with enterprise strategy improves resource allocation, reduces unnecessary duplication of effort, and increases the likelihood that AI investments generate measurable organizational value over time.

Once a legitimate business need has been established and organizational value has been demonstrated, AI adoption transitions from identifying why artificial intelligence should be implemented to determining how the organization will support its adoption. Within the AI Workflow Lifecycle Framework, Business Need is intentionally positioned outside the lifecycle because it acts as the organizational trigger for AI adoption rather than an operational phase of AI implementation. The next organizational capability, Executive Sponsorship, provides the strategic commitment, organizational authority, and long-term investment necessary to transform a validated business need into an operational organizational capability.

V. Proposed AI Workflow Lifecycle Framework (AWLF)

Existing frameworks such as the NIST Cybersecurity Framework (CSF) 2.0 and the NIST Artificial Intelligence Risk Management Framework (AI RMF) provide organizational guidance for managing cybersecurity and AI risk. However, organizations implementing enterprise AI workflows must coordinate governance, security, executive leadership, and lifecycle management simultaneously. To address this need, this paper proposes the AI Workflow Lifecycle Framework, a conceptual organizational framework for managing AI workflows as organizational capabilities rather than software projects.

As illustrated in Figure 1, the framework begins with a validated business need that initiates the AI adoption process. Executive sponsorship provides organizational commitment, governance establishes oversight and accountability, and security continuously surrounds the framework as an enduring organizational function. At the center, the AI workflow lifecycle consists of development, operations and maintenance, and continuous improvement.

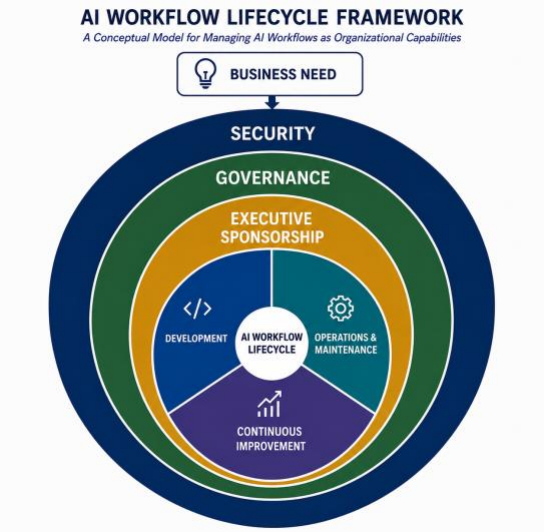

Figure 1. Proposed AI Workflow Lifecycle Framework

The AI Workflow Lifecycle Framework illustrates the organizational capabilities required to successfully implement and sustain enterprise AI workflows. A validated business need initiates the framework, executive sponsorship establishes organizational commitment and ensures necessary resources throughout the lifecycle; without this support finances may start to fizzle out or be reallocated after development, maintenance and improvement will also need a budget if the organization wants the adoption to be successful. Governance provides strategic direction and oversight, and security continuously surrounds the framework as a persistent organizational function. At the center, the AI workflow lifecycle consists of three operational phases: development, operations and maintenance, and continuous improvement.

The beginning of the AI workflow implementation process must always be based upon business need. If a Chief Executive Officer (CEO) or board member gets overzealous, because they heard about another company using AI to automate business processes that enthusiasm should be channeled and redirected. AI workflows in the enterprise capacity are not a new software project that will suddenly generate an organization more revenue and improve productivity without any overhead. They must be planned through a legitimate, structured process as would any other organizational change. The first sign that an organization could benefit from the implementation of AI workflows is real, tangible business need. If the workflow could solve a legitimate problem within the enterprise, one that executives would be comfortable pitching to shareholders then the organization should continue with the design and planning process. It will immensely benefit the organization to thoroughly ensure that there is significant organizational value, instead of finding out upon implementation that the decision was driven largely on emotion and short-term market trends instead of a strategic business decision that will place the organization in a favorable position in the market. If the consensus is that implementing a particular AI workflow will be build long-term shareholder value while fulfilling a legitimate problem or gap in the business, then an organization will likely have success through the next steps of this framework, creating an AI workflow lifecycle, not a short-term project.

After the business need successfully triggers the first stage of the lifecycle an organization can begin the planning phase. The planning phase is where you begin designing the AI Workflow in accordance with executive sponsorship. In this framework there will not be an explicitly outlined design process, as the AI Workflow Lifecyle Framework (AWLF) focuses on defining operational procedures and steps, instead of acting as an implementation playbook.

Executive sponsorship is critical to elevating AI workflows from a software project to an organizational capability. Sponsorship ensures that the sentiment is shared amongst the executive leadership team of an organization and that teams or individuals involved in the planning, design, and implementation process have adequate resources and support. This also ensures that support won't decline after development, ensuring that the AI workflow sees all stages of the lifecycle stated in this framework. With executive sponsorship an organization can ensure that funding is properly secured not only for initial development but through every stage of the continuous lifecycle. AI workflows are not systems that can be left to run autonomously without any maintenance or further improvement. Through this sponsorship a team can ensure this project is elevated from a business solution to an enterprise-wide initiative. Once the momentum from the organization-wide initiative has "kickstarted" the process it is crucial to determine strategic direction. Every piece of the process should be structured to ensure success in later

stages. This executive support will also help the initiative obtain the utmost levels of coordination amongst departments, as it is no longer a department project, it is a company-wide initiative. Combined, this should remove organizational roadblocks and barriers, creating the best chance of success for the project. The benefits created by executive sponsorship should accelerate the process and improving the odds that the project doesn't become abandoned as a proof of concept.

Throughout the executive sponsorship process and every step in the framework, governance should be considered. In the diagram, governance is an outer ring encompassing the other organizational functions, as it should be thought of as a continuous process. Even beyond the lifecycle of the workflow, there must be significant oversight and governing of the AI workflow systems. They must always be considered critical systems, even after End-of-Life (EOL), as they handled company data. Governance surrounds every decision made regarding workflows, as policies must be made and executed regarding these systems throughout every stage of the lifecycle. The incorporation of AI workflows into an enterprise will require organizational roles to be defined, whether it be how a role interacts with, designs, or maintains them. Beyond defining roles, accountability and responsibility must also be designated. Governance will also help determine a risk appetite the organization is comfortable with. Governance is clearly a lifecycle long commitment, as although many of these objectives will be completed in the planning process, oversight is required to ensure that the workflow does not exceed the accepted risk appetite or that responsible parties are taking accountability. This framework utilizes the governance function in the same capacity and philosophy as the National Institute of Standards and Technology (NIST) Cybersecurity Framework 2.0 Govern or NIST AI Risk Management Framework functions.

Although Governance is incorporated into every piece of the lifecycle, security should be involved in every phase or function including within every governance decision. AI workflows should be completely secure-by-design; meaning that every decision should be made with security in mind and it should not become actionable until security concerns have been properly considered. Artificial Intelligence in itself introduces new attack surfaces and vulnerabilities that require security professionals and executive leadership to take a different approach than with traditional systems. Incorporating AI into workflows, giving them access to enterprise software and confidential data heightens the security risk for the company, thus eating into that risk appetite set under the governance function. Confining security to development leaves gaps in planning and design and omits much of the process of implementing operational security controls. It is important to ensure security is involved from the beginning of the process. Throughout the executive sponsorship process it is imperative to draw support both organizationally and financially for security throughout the entire lifecycle. Security begins here, is vital throughout the development process, should be considered continuously throughout operation and is necessary for successful improvement. As technology and business processes continue to evolve, the AI workflow will continue to change and grow. It is important that as a secure-by-design system security is not omitted here, it is increasingly relevant throughout the entire AI workflow evolution process.

Beyond these guiding functions, the AI workflow must advance from a proof-of-concept to an actual living system. Now that the organization has made the necessary preemptive decisions to build the workflow, the project can begin the development phase of the lifecycle. Development should be structured, ensuring the workflow is secure-by-design. Developers should ensure code and databases are secure from the ground up. Security shouldn't be an afterthought. They should also ensure the longevity of the workflow and that it is possible for themselves or future developers to improve or innovate upon in the future.

Just as governance and security are continual, the work after development doesn't end there. Although AI workflows are supposed to act more autonomously than chatbots in regards to business processes, they still require consistent maintenance. AI workflows cannot be pushed from development to production and left alone. They need staff to conduct maintenance in order for AI workflows to operate reliably. This will ensure they maintain the necessary level of performance that contributes business value to the organization.

As an organization continues to maintain AI workflows, they should also ensure they are continuously improving the workflows, so they bring more value to the organization and continue to give the company a competitive edge. As technology continues to evolve business requirements and risks will evolve with it. This is why it is crucial to continuously develop and improve AI workflows as AI capabilities improve. Models will continue to evolve, and workflows should evolve with it.

Many of these operational processes overlap with one another, when combined as a unified framework they establish an approach to enterprise AI adoption that prioritizes treating AI workflows as a unified lifecycle. This will largely benefit the success of enterprise AI workflows as it allows workflows to be built and managed in a structured process where each step strengthens and interacts with the other instead of a solitary process where each step functions independently of another, creating systematic inefficiency. The AI Workflow Lifecycle Framework views AI workflows as a process that requires continual organizational commitment rather than as a discrete software

project. The AWLF emphasizes that enterprises should make AI workflows an organization-wide initiative. Executive leadership, governance, security, and continuous improvement are necessary to successful enterprise AI workflow adoption.

VI. Conclusion

As AI evolves from isolated tools to structured organizational capabilities, enterprises require governance models that extend beyond risk management and software development frameworks. The AWLF provides a foundation for this necessary transition by highlighting the need for identification of a verified business need and executive commitment. Combined with overall organizational alignment and continual security and governance functions, the organization will experience a structured, successful adoption of AI. This framework suggests that AI workflow implementation will be more successful if they are implemented with security, clear governance policies, and lifecycle management. However, this is primarily conceptual; future work would benefit the AWLF by evaluating it across organizations that vary in size, industry, and mission.

References

Bommasani, R., Hudson, D. A., Adeli, E., Altman, R., Arora, S., von Arx, S., Bernstein, M. S., Bohg, J., Bosselut, A., Brunskill, E., Brynjolfsson, E., Buch, S., Card, D., Castellon, R., Chatterji, N., Chen, A., Creel, K., Davis, J. Q., Demszky, D., ... Liang, P. (2021). *On the opportunities and risks of foundation models*. Stanford Center for Research on Foundation Models. https://arxiv.org/abs/2108.07258

Brown, T. B., Mann, B., Ryder, N., Subbiah, M., Kaplan, J., Dhariwal, P., Neelakantan, A., Shyam, P., Sastry, G., Askell, A., Agarwal, S., Herbert-Voss, A., Krueger, G., Henighan, T., Child, R., Ramesh, A., Ziegler, D. M., Wu, J.,

Winter, C., ... Amodei, D. (2020). Language models are few-shot learners. *Advances in Neural Information Processing Systems, 33*, 1877–1901.

Coyne, M., Larson, J., Shieh, J., & Yeon, H. (2019, December 5). *Winning in automation requires a focus on humans*. McKinsey & Company. https://www.mckinsey.com/capabilities/tech-and-ai/our-insights/winning-in-automation-requires-a-focus-on-humans

Deniz, B. K., Gnanasambandam, C., Harrysson, M., Hussin, A., & Srivastava, S. (2023, June 27). *Unleashing developer productivity with generative AI*. McKinsey & Company. https://www.mckinsey.com/capabilities/tech-and-ai/our-insights/unleashing-developer-productivity-with-generative-ai

Manyika, J., & Sneader, K. (2018, June 1). *AI, automation, and the future of work: Ten things to solve for*. McKinsey & Company. https://www.mckinsey.com/featured-insights/future-of-work/ai-automation-and-the-future-of-work-ten-things-to-solve-for

McKinsey & Company. (2025). *The state of AI: How organizations are rewiring to capture value*. https://www.mckinsey.com/capabilities/quantumblack/our-insights/the-state-of-ai

National Institute of Standards and Technology. (2023). *Artificial intelligence risk management framework (AI RMF 1.0)* (NIST AI 100-1). U.S. Department of Commerce. https://doi.org/10.6028/NIST.AI.100-1

National Institute of Standards and Technology. (2024). *Artificial intelligence risk management framework: Generative artificial intelligence profile* (NIST AI 600-1). U.S. Department of Commerce. https://doi.org/10.6028/NIST.AI.600-1

Peng, S., Kalliamvakou, E., Cihon, P., & Demirer, M. (2023). *The impact of AI on developer productivity: Evidence from GitHub Copilot*. Microsoft Research. https://www.microsoft.com/en-us/research/publication/the-impact-of-ai-on-developer-productivity-evidence-from-github-copilot/

Pietrzak, M. (2025). A trillion dollars race—how ChatGPT affects stock prices. *Future Business Journal, 11*(1), Article 50. https://doi.org/10.1186/s43093-025-00470-5

Russell, S. J., & Norvig, P. (2021). *Artificial intelligence: A modern approach* (4th ed.). Pearson.

Vaswani, A., Shazeer, N., Parmar, N., Uszkoreit, J., Jones, L., Gomez, A. N., Kaiser, Ł., & Polosukhin, I. (2017). Attention is all you need. *Advances in Neural Information Processing Systems, 30*.