



WHITE PAPER

A FRAMEWORK FOR MANAGING A COMPANY WIDE AI PROGRAM

Dr. Brian Scott Glassman



- • •
- • •
- • •
- • •
- • •
- • •
- • •

Article Summary

This article presents a framework that outlines the essential areas, processes, and methods for managing an AI program to maximize its value creation for corporations. The visual framework is divided into three sections: stakeholders, AI leadership, and the AI portfolio, which includes a pipeline of use cases, proof-of-concept projects, pilot projects, and company-wide rollouts. This article is the first in a series on AI program management, with subsequent articles delving into each aspect of managing a corporate AI program.

Written for executives, VPs, AI practitioners, and stakeholders, this article provides a valuable framework for managing AI initiatives. By focusing on the management aspects of AI programs rather than the technologies themselves, the content of this article remains relevant and applicable for the foreseeable future. Readers are encouraged to follow the author for future articles that will provide in-depth insights on various aspects of managing AI programs within organizations.

I. The Immediate Need to Establish a Corporate AI Program

Throughout his experience in management and technology consulting, particularly in the fields of Artificial Intelligence, Machine Learning, and Generative AI, the author has closely observed the strategic initiatives undertaken by corporations. These endeavors have exhibited diversity in their organizational frameworks, leadership approaches, and deployment methodologies. However, due to their newness many of these initiatives have yet to yield significant results.

In the past two years, Generative AI has demonstrated its transformative potential, with industry leaders such as OpenAI, Anthropic, Mistral, and NVIDIA at the forefront. This disruptive technology has rapidly revolutionized work practices across sectors like creative writing, marketing, publishing, and coding. As this evolution continues, even more powerful and disruptive AI technologies are anticipated to emerge.

Thankfully, this perspective is shared by the broader business community, which is eager to leverage AI applications within their organizations. The unprecedented pace at which AI is evolving and reshaping business strategies parallels the transformative impact of the internet's emergence, which rendered numerous once-dominant corporations obsolete across industries such as retail, media, telecommunications, and advertising. Companies are on the verge of another significant disruption, with leaders like Google, Facebook, and Microsoft actively understanding how to adapt their business models.


- Companies lacking AI expertise and proximity to those technologies will need to acquire similar AI capabilities to effectively integrate it into their operations and ensure their survival. Now is the opportune moment to establish AI leadership, formulate strategic frameworks, and implement organizational structures to confront this challenge and capitalize on the associated opportunities.
- This article underscores the imperative for proactive corporate measures, with the aim of guiding corporations towards a comprehensive understanding of AI deployment and providing them with strategic exposure to key AI technologies that can impact their business before it is too late.

II. Benefits of an Organized AI Program

Organizing AI into a formal program with stakeholders, AI leadership, and a pipeline of projects is essential for four main reasons. Firstly, it openly acknowledges and demonstrates visible commitment to AI. Secondly, it provides a structured approach that enhances the likelihood of success for each AI deployment. While a piecemeal approach may suffice for basic AI deployments like ChatGPT or MS Copilot, more complex endeavors require a supportive ecosystem of feedback and support to avoid floundering or failure. The real power of AI lies in leveraging the growing variety of AI applications to address your corporation's specific strategic needs effectively, and this framework provides that support. Furthermore, having a formal program helps in formalizing a budget and provides visibility into the progress of each specific AI initiative. Lastly, being able to measure and track the cost versus impact of multiple AI initiatives from conception to deployment is something the CEO and board will greatly appreciate, and further fuel the momentum behind AI's.


III. A Framework for Managing a Corporate AI Program

This article introduces a visual framework for managing corporate AI programs. This framework integrates AI leadership, stakeholder management, and AI portfolio management. Each area of the framework will be discussed, including how it can be implemented and general rules of thumb for managing that area. The following articles will dive into details for each area, follow the author on medium.com for future publications.



-
-
-
-
-
-
-
-
-
-

The Glassman framework is formed by three key layers. At the top are the stakeholders, followed by the leadership tier known as AI leadership and management. Anchoring the structure is the portfolio of AI projects, spanning from targeted use cases to expansive company-wide initiatives. Let us examine the uppermost layer of AI stakeholders.



Section 1. AI Stakeholders

In the AI deployment landscape, stakeholders play key roles by supporting initiatives and providing strategic insights aligned with corporate goals. Most importantly, stakeholders provide feedback on AI initiatives, help in removing roadblocks, and work to evangelize the new AI to help in its adoption.

Now, there are two main groups of stakeholders: internal and external stakeholders. Let's start with internal stakeholders. These stakeholders should consist of representatives from the C-suite, VPs, department heads, technology, and project management. However, when selecting internal stakeholders, it is important to carefully ensure diverse perspectives and consider the stakeholders' abilities. External stakeholders must also be carefully chosen and vetted to ensure there are no conflicts of interest. For both internal and external stakeholders, it is crucial that they provide constructive feedback and demonstrate a genuine interest in advocating for the organization's AI initiatives, without being roadblocks themselves.

Numerous best practices exist for selecting, managing, and communicating with stakeholders. However, to maintain the article's focus, this information will be provided in a separate piece. Follow the author to stay informed of future insights.



Section 2.0. AI Leadership: The Key to Success

The second layer of the framework encompasses AI management, which includes the Head of AI, research initiatives, the establishment and maintenance of guiding ethics and policies, AI risk management, and ensuring AI compliance with both industry-specific regulatory laws and future AI policies at national and international levels.

- • •
- • •
- • •
- • •
- • •
- • •
- • •


At the helm of the AI department is the Head of AI, who plays a crucial role in the ongoing success of AI initiatives. Companies should resist the urge to hire internally or within their industry and instead seek AI experts, particularly those specializing in generative AI. The knowledge required to excel in AI/ML and generative AI is applicable across industries and is advancing rapidly. Therefore, the Head of AI should be a knowledge leader in AI rather than solely an industry expert. They should possess the ability to innovate across the organization, gather necessary resources, and act as a change agent to reduce resistance to new innovations that will transform work processes and customer interactions.

Equally critical to the success of AI initiatives is the research component. AI research is vital due to the unprecedented speed of technological advancements. The AI department must track multiple areas of research to identify new developments that can aid in company-specific AI initiatives, reduce the associated risks, and improve the AI software's reliability. While academic research was traditionally the primary source of AI knowledge, AI's pace renders many academic articles outdated quickly. Consequently, more informal research methods should be adopted. Additionally, research should include evaluating vendor AI solutions with a mind toward their pace of development as today's solutions may greatly improve in the near future.

Ethics and policies for AI are a vital pillar in the AI department. It starts with understanding the risks that the company can face with AI and setting up policies and procedures that can mitigate these risks. For example, releasing a generative AI product that can be tricked into providing restricted or sensitive information is an issue that should be addressed in policies limiting who is permitted to access the AI. Policies and ethics are a complex topic, and it will be addressed in a follow-up article, follow the author for more information.

Risk management is an area that needs to be researched and understood through investigation and testing of the company specific AI software programs. As AI is a black box, extensive testing using edge cases is very necessary. Furthermore, policies limiting the deployment of AI are needed until the range of behaviors can be well understood. Risk management also extends to the use cases. Here the AI should be limited to a select set of use cases, and then later after extensive testing can be opened up to additional use cases. Additionally, training users on the limitations and risks associated with the AI system is a vital task.

Finally, regulations specific to the industry, country, and business area need to be applied to the specific AI application. This requires both knowledge of the regulations and an understanding of how the AI applications could possibly circumvent those regulations. If the AI is for internal use, strict training needs to be provided on how to and not to use the AI application and why. If the AI is for external use by customers or other parties, SW based safeguards need to be put in place and tested with edge cases along with legal disclosures.



Section 3.0. The AI Portfolio: Deploying Company-Specific AI Initiatives

The AI portfolio is a strategic approach comprising four distinct phases designed to effectively integrate artificial intelligence into an organization. The first phase involves discovering and defining the use cases, followed by the second phase, which focuses on creating and testing the AI through pilot projects AKA proof of concepts projects. The third phase consists of conducting test deployments, and finally, the fourth phase involves implementing company-wide rollouts of AI initiatives.

It is crucial for companies to resist the temptation to skip from the first phase directly to the fourth phase, as this can lead to significant risks, particularly when dealing with AI technology. AI systems can be considered “black boxes” due to their inherent complexity, and limited ability to see their inner workings. Thus one must have a staged testing plan to insure reliable outputs. Furthermore, organizations require sufficient time to familiarize themselves with the technology and establish appropriate rules and limitations based on the results of their testing.

A cautionary example is Google’s February 2024 Gemini release, where inputs far upstream during the training of the large language model had unforeseen consequences on the AI’s outputs, ultimately damaging the company’s reputation. This incident highlights the importance of adhering to standard processes and thorough testing via the second and third phase of this framework.

The four phases of the AI portfolio are designed to serve a multitude of objectives, all of which contribute to the successful integration of AI technology within an organization. This phased approach helps to mitigate the risks associated with AI implementation, while simultaneously providing the company with valuable opportunities to learn how to effectively operate and monetize these technologies. Additionally, the phased approach allows for the adjustment of software architectures, the development of best practices for AI use cases, the creation of governance policies, the creation of measurement for monitoring the AI’s impact, and the formulation of change management plans.

- • • Perhaps most crucially, the four phases enable the prioritization of initiatives, allowing the allocation of resources to the most promising AI initiative and the termination of less viable ones prior to a company-wide rollout. This strategic approach ensures that the organization invests their resources wisely and minimizes the risks of investing in projects that may not yield the desired results. Moving forward, let us delve into the first phase of the AI portfolio: Creating the use cases and developing their associated business cases.



3.1. Phase 1. Use Case Exploration and Validation

The initial phase of the AI portfolio level focuses on exploring and developing use cases, along with creating their corresponding business cases. A use case for AI is defined as a specific set of activities performed by an employee, partner, or customer in which they interact with the AI to achieve a particular goal. It is important to note that simply applying ChatGPT to a company does not constitute a use case; however, being more specific by utilizing it within the marketing department to enhance the quality of marketing collateral does qualify as one. Use cases can be either internal or external-facing and can range from highly niche to very general in nature. Regardless of their scope, use cases must have measurable outcomes. In the aforementioned example, the assessed average quality of the marketing collateral before and after the implementation of ChatGPT would serve as a quantifiable outcome. Moreover, use cases can have multiple measurable outcomes. With a well-defined use case and its predicted measurable outcomes, a business analyst can develop a business case. For instance, if the projected quality of marketing materials increases from a rating of 5 out of 10 to an 8 of 10 after the implementation of AI, and the project time required to produce the finished collateral decreases by 30%, a strong business case can be made in hours saved and improved brand quality.

3.2. Phase 1. Technology Research

To fully develop business cases, insights from the AI program's team are necessary to determine the technology requirements. It is crucial to consider that the rapid pace of AI developments may render previously infeasible AI use cases now feasible. Therefore, as new technology solutions emerge, the use cases in the queue should be updated if they are positively impacted.

- • •
- • •
- • •
- • •
- • •
- • •
- • •

AI analysts and business analysts must collaborate to create business cases for each use case. The AI team will provide valuable information such as potential risks, required development resources, investment amounts, make-or-buy options, and timelines for development. Additionally, the AI team member will broadly identify the AI technologies needed for the use case and provide a rough estimate of the usage costs.

3.3. Phase 1. Business Cases Development

Business cases are more comprehensive than use cases, as they include potential risks, required development resources, investment amounts, make-or-buy options, idealized timelines for development, and return on investment (ROI). Crucially, AI business cases provide an estimated range of monetary impact on the business. It is important to note that this estimate does not need to be precise, as the subjective and objective outcomes, based on key performance indicators (KPIs), will be better understood and further refined as the AI initiative progresses through the different phases. When presenting business cases to stakeholders, it is essential to ensure that they are easy to read and omit unnecessary technical details.

3.4. Phase 1. Identifying Opportunities for AI Use Cases

Discovering use cases across a large organization can be a challenging task. However, several methods can be employed to identify use cases effectively. Some approaches involve examining AI technology to determine common applications, while others are more complex but result in strategically aligned AI use cases with greater impact. A detailed article addressing this topic is forthcoming and is one of the author's areas of expertise, follow for more information.

Now, the process of finding use cases, known as opportunity mapping, can uncover numerous high-value AI opportunities when executed effectively. Conversely, a poorly performed exercise may only identify common use cases. Investing time and effort in finding AI use cases by leveraging stakeholders and internal company resources can greatly increase the quality of the resulting use cases and ensure alignment with the company's strategic goals. Hence, shortcutting the opportunity identification process is not recommended.


3.5. Phase 1. Stakeholder Engagement

It is crucial to provide stakeholders with regular updates on the progress of uncovering AI use cases and to regularly deliver formal presentations showcasing the identified use cases. These formal presentations allow the stakeholders to thoroughly examine the use cases, offer valuable insights, and guide the AI team toward relevant resources within the organization. Furthermore, these presentations enable stakeholders to contribute additional ideas for use cases, enhancing the overall quality and breadth of the AI portfolio.

- • •
- • •
- • •
- • •
- • •
- • •
- • •

Under the guidance of the Head of AI, stakeholders will contribute thoughts on which use cases from the portfolio should advance to the pilot phase of testing. The selection of pilot projects should be a collaborative effort, taking into account a balanced mix of factors such as AI application risk, time-to-implement, required resources, potential revenue impact, and the likelihood of internal adoption. It is important to ensure that the pilot portfolio is not overly skewed in any particular direction. For example, selecting all low-risk, easy-to-implement projects or conversely focusing on high-risk, large-scale initiatives can jeopardize the success of the entire AI program. Striking the right balance of AI initiatives in the AI portfolio is essential to mitigate risks and maximize the potential for success.

By engaging stakeholders in the selection of AI initiatives and carefully considering the various factors that influence the viability of each use case, organizations can create a well-rounded use case portfolio that aligns with their strategic objectives and sets the stage for a successful AI implementation.



4. Phase 2. Pilot Projects

The second phase of the AI initiatives pipeline involves the execution of pilot projects, also referred to as proof-of-concept (POC) projects. During this phase, a structured project management framework is established to oversee and guide the progress of these projects. The quantity of POC projects undertaken is contingent upon the organization's capacity and available resources. When implementing POC projects, it is crucial to define clear learning objectives, establish technology-related goals, and outline key performance indicators to measure success and gather valuable insights.

4.1 Phase 2. Learning Goals of the POC

It is generally advisable to conduct pilot projects separately from the company's core operations or customized workflows, ensuring that their outputs are used solely for evaluation purposes. Learning goals should encompass comprehensive assessments of company and customer outcomes across diverse scenarios. This approach may necessitate the replication of work, such as reproducing a marketing brochure using a large language model (LLM) AI and conducting a comparative analysis against the previously created brochure. The results of using the AI should facilitate an apples-to-apples comparison.

- • •
- • •
- • •
- • •
- • •
- • •
- • •

Another illustrative example involves employing an AI customer service agent, powered by OpenAI's ChatGPT, to assist customers. In this case, internal employees would assume the role of simulated customers, engaging with and testing the AI service agent. The outcomes of these interactions would be formally evaluated and rated through a structured feedback survey. Comprehensive learning plans, incorporating a wide array of agent-customer interactions, should be executed, simulating various scenarios and comparing the results to those of existing systems. The feedback gathered from simulated customers would be relayed to the AI team, facilitating the refinement of business cases.

Moreover, learning plans should incorporate an assessment of potential risks and edge cases. For instance, consider a situation where an AI chat agent provides inaccurate information, potentially exposing the company to legal liability. By proactively identifying and addressing such risks, organizations can mitigate potential adverse consequences and ensure the responsible deployment of AI technologies.

4.2 Phase 2. Technology Learning Goals

The landscape of AI technologies is rapidly evolving, with a proliferation of options emerging on a daily basis. Large language models (LLMs), for instance, are available in various sizes, capabilities, and price points. The POC technology team bears the responsibility of evaluating an array of AI technologies and frameworks against a set of well-defined metrics to identify the most suitable solution for a particular use case. Understanding the impact of the AI technologies employed in the pilot project on the resulting outcomes is of paramount importance.

To illustrate this point, let us revisit the AI customer service agent example mentioned earlier. An LLM such as Llama 2 70B may demonstrate the ability to handle straightforward customer interactions effectively. However, it may struggle to address more intricate or complex inquiries. In contrast, ChatGPT 4 may possess the capacity to navigate these edge cases with greater proficiency but at a higher cost. The technology team must carefully consider the trade-offs between cost, performance, and risk when making decisions regarding the selection of AI technologies.

It is crucial for the project management and technology teams to maintain close collaboration and open lines of communication throughout the POC process. By working in tandem, they can develop a solution that optimizes various trade-offs and aligns with the organization's overarching objectives. This synergistic approach ensures that the chosen AI technologies not only meet the technical requirements but also contribute to the realization of desired business outcomes while mitigating potential risks.


4.3 Phase 2. Decision Point

Throughout the POC process, a critical decision point is reached when enough learning has been accumulated to determine the future of the AI initiative. The decision involves three potential paths: 1) placing the POC on hold until future technology enablers are available, 2) advancing the POC to the test deployment phase, or 3) terminating the POC.

If a use case shows promise but the technology is too expensive or complex, the POC should be placed on hold. If the POC meets the requirements in the business cases, it should undergo further de-risking and evaluation through test deployments. Finally, if the POC fails to meet holistic expectations, such as yielding only incremental productivity gains or delivering inconsistent results, it should be terminated.

Terminating a POC does not necessarily mean the use cases were flawed, but rather that the current technology and business trade-offs were not optimal. Terminated use cases can be revisited as AI technologies advance and business priorities shift. Given limited internal resources, POCs must compete for resources and prioritization via their results.

Finally, maintaining a strategic and adaptable approach to AI initiatives is crucial for maximizing return on investment and ensuring alignment with evolving business objectives. In the future, vendors may introduce a superior solution that renders AIs built in-house less advantageous from a resource allocation perspective.



5. Phase 3. Test Deployments

In the third phase of the AI initiatives pipeline, the project progresses to a limited test deployment. This phase involves the integration of evangelists, project management, stakeholders, governance, and formal key performance indicators (KPIs). Each of these components plays a crucial role in providing feedback, monitoring performance, ensuring compliance, and gathering metrics to build a compelling case for advancing the project to subsequent phases. Governance and KPIs are introduced in this phase and warrant further explanation.


- • •
- • •
- • •
- • •
- • •
- • •
- • •

AI solution governance encompasses the designation of an application owner who are responsible for overseeing the solution and making swift decisions, particularly in high-risk situations. This owner need not be a member of the AI team. Governance also extends to the development and adaptation of legal frameworks and policies as the testing the AI deployment progresses. If the AI application is internally focused, training on governance and policies must be provided to the relevant personnel.

KPIs should be collected through both automated systems and informal means to assess the performance of the AI application. For instance, an internal AI Chat deployment should automatically track employee usage patterns, while the project management team conducts one-on-one interviews to gather pointed user feedback. These feedback mechanisms are essential for making necessary adjustments to boost the internal adoption rate of that AI technology.

Further, a combination of technological and organizational changes must be implemented and closely monitored. The presence of change agents and evangelists on the team significantly increases the likelihood of successful adoption and aids in effective planning. Incentivizing these individuals is also crucial to further bolster adoption and usage rates with the test deployment.

Finally, due to internal capacity constraints, it may only be feasible to accommodate one or two deployments concurrently during the test deployment phase. Portfolio decisions made by the AI leadership, guided by stakeholder input, should be performed for the selection of AI initiatives placed in this phase. Throughout this phase, a substantial amount of learning, adjustments, and structured processes should be established to maximize the adoption of the AI projects while minimizing associated risks.



6. Phase 4. The Company Wide Rollout of the AI Initiatives

The company-wide rollout of AI initiatives is a critical objective for an AI program, requiring several essential intermediary steps to ensure successful adoption. These four-phases emphasize the importance of fostering the adoption of the AI initiative through the implementation of change management strategies, comprehensive communication plans, and incentive structures. Kotter's 8-Step Change Model will serve as an effective framework for this process.

To facilitate a smooth and effective implementation, it is crucial to provide comprehensive communication plans and training to all individuals involved in the AI rollout. Incentive plans should be designed and implemented, particularly for evangelists and lead users, to encourage their active participation and support. Moreover, lessons learned from previous governance policies should be incorporated into the rollout plans.

Here it is essential to set and communicate the goals and KPIs of that particular AI initiative. Those involved in the corporate rollout should have a clear understanding of the anticipated benefits of that AI initiative and support the AI team in further providing feedback via formal mechanisms. Providing the organization with a clear understanding of the anticipated benefits and outcomes is crucial for fostering engagement and support.

From a risk management perspective, each rollout should be closely monitored. It is important to recognize that AI systems that perform well during test deployments may exhibit changes in behavior due to various factors. If the AI starts to behave in an unusual or unexpected manner, it is crucial to be prepared to pause the rollout and re-assess the technology. This proactive approach ensures that any potential risks or issues are identified and addressed promptly, maintaining the integrity and effectiveness of the AI initiatives.

By adhering to these practices and employing a structured approach to the company-wide rollout of AI systems, organizations can maximize the potential benefits of their AI initiatives while minimizing risks and ensuring successful adoption across the enterprise.

Conclusions

In conclusion, this article presents an initial overview of a framework designed to effectively manage a corporate-wide AI program. The framework outlined in this article is intended to provide a structured approach to implementing and overseeing AI initiatives across an organization.

If the readers need further clarification or would like a more detailed explanation of any specific aspect of the framework, please do not hesitate to reach out to the author. Your feedback and inquiries will also serve as valuable guidance for his future writing endeavors. Please do connect with the author on LinkedIn, and follow on Medium.com, and like if this article was valuable.