AI Best Practices for Business Decision Makers and Practitioners

Contents

Introduction

Artificial intelligence (AI) is a branch of computer science that uses algorithms (programmatic instructions) to process inputs and deliver outputs that mimic human abilities such as cognition, natural language understanding, speech and image understanding, common sense reasoning, knowledge representation. Machine learning (ML) and deep learning (DL), the subfields of AI driving innovation and headlines, have introduced improvements to pattern recognition thereby supercharging tasks such as classification, prediction, and clustering.

With companies in every industry and geography leveraging AI, or at least wanting to, it has never been more important for technical practitioners as well as non-technical decision makers to understand how AI can benefit their business as well as the associated risks of implementing AI into the enterprise. It is critical for key stakeholders to articulate the business value of utilizing AI for solving their business problems, and to understand the associated costs and benefits of deploying AI-infused applications, the time to value of implementing AI and what success will look like over time.

The rampant confusion and conflation surrounding AI terminology is a major barrier to entry for most enterprises, no matter how big or small. For example, it is estimated that 40% of companies claiming to be "AI startups" show absolutely no evidence that the technology is material to the execution of their value proposition. At the same time, in 2020, companies are just beginning to understand the risks posed by AI systems that do not function in ways that are expected or can be explained.

For example, in August 2019, a large tech firm partnered with a financial services firm to co-launch a credit card that was quickly discovered to offer women exorbitantly lower credit limits than men, leading the tech company's co-founder to describe the situation as having used a misogynistic algorithm. Two months later, a study published in the American Association for the Advancement of Science revealed that an algorithm being leveraged by a large healthcare organization to optimize the outcome of hospital visits was providing comparable risk scores to white and Black patients, despite the Black patients being significantly sicker, leading to their receiving disproportionately insufficient care.

> According to IDC, global AI technology spend was estimated at approximately $50B in 2020 and is expected to double to $110B by 2024.

Amid that backdrop, the global AI technology spend was estimated at approximately $50B in 2020 and is expected to double to $110B by 2024, according to IDC. Across industries, companies are accelerating efforts to integrate various machine learning and deep learning frameworks while adhering to best practices and installing proper safeguards.

In this guide, the CompTIA Artificial Intelligence Advisory Council leverages the diverse and expansive domain expertise of its members to provide curated insights on pain points and best practices associated with infusing AI solutions within an organization. To provide effective guidance, the information presented here is organized into best practices for two personas that are frequently involved in the adoption and integration of AI technologies within any organization: the practitioner and the decision maker.

Practitioners work directly with data and technology and support decision makers' strategic goals to achieve their desired business outcomes. They tend to have job titles such as data scientist, data engineer, data analyst, or machine learning architect. Practitioners are concerned with the theory, mechanics, and practical applications behind the underlying algorithmic and mathematical approaches. They translate a business problem into a mathematical problem that can take data produced by business applications as inputs and produce outputs in the form of actionable insights. They work closely with IT to build the tools and infrastructure needed to deliver scalable software solutions and data-driven optimizations.

Decision makers are often focused more on the macro-level goals and challenges of the business. They explore ways in which modern technologies can assist them in addressing current and future challenges in the most cost-effective manner that offers a generous return on investments. They tend to have job titles such as founder, CXO, CRO, line of business (LOB) leader, chief data officer, vice president of business intelligence, or vice president/director of engineering/IT. They are often tasked with performing cost-benefit analyses and data visualizations that allow non-technical stakeholders to quickly understand the business case for implementing AI-driven solutions. They are also keen to ensure compliance and data standards utilized across the business.

Of course, the lines between the two personas can blur depending on the size or scale of an organization. For each of the personas, this guide enumerates and expands upon the key pain points and best practices that are likely to be encountered and relevant along a business's AI journey.

- *Pain points* are specific tasks or requirements that consistently represent difficulties for a particular type of stakeholder.
- *Best practices* are the procedures that are generally accepted, particularly among experts, as being the most effective, based on currently available research.

While every company and use case is different, these guiding principles should maximize the efficiency and efficacy of your AI initiatives. As modern technologies and applications are researched and developed, pain points may dissolve while best practices continue to evolve. This report reflects the current understanding and landscape of available AI approaches and should be interpreted within that context.

Part 1: Decision Makers

Infusing artificial intelligence (AI) into businesses is challenging. Decision makers are often tasked with solving a business pain point with state-of-the-art technology solutions. Deploying solutions, especially those that involve AI, come with their own pros and cons. One must worry about collecting enough data to train those AI models and ensure that those models are fair and unbiased, transparent, explainable, accurate, continuously learning, fresh, secure, and robust against adversarial attacks. Business decision makers also must ensure that there are sufficient skills in the organization to maintain AI models that are developed as these are a new breed of applications. In this section, we describe best practices for some of the challenges decision makers face when deploying AI-infused applications.

Pain Point	Description
Shiny Object Syndrome	There is ever-growing buzz around AI and an incessant coverage of how AI can solve all problems. Due to this increased market and peer pressure, businesses are looking to jump on the AI bandwagon without understanding the cost and implications associated with AI or AI-infused solutions. Technology is rapidly changing so having a roadmap is critical to avoid the high cost of project churn, both in application and infrastructure. How can decision makers infuse AI into their businesses while providing meaningful and tangible value?
Business Use Cases	Wanting to use AI is one thing but identifying specific business processes and use cases that will benefit from AI infusion is an entirely different ball game. Decision makers are faced with the challenge of identifying and prioritizing key business processes and use cases that will benefit from AI and provide a good return on investment (ROI).
Data Discovery and Management	AI needs data. In essence, data is the fuel for training and inferencing. Decision makers are required to work with the leaders of various departments in their organization to identify existing data silos and overall quality and health of the data. They should strategize and recommend ways to streamline and manage the data. This will help the business to tap into the data to discover trends and gather new insights.
Model Discovery and Management	A fair amount of experimentation is required to find a reasonable AI/ML model. Reasonable is defined by the acceptable accuracy of the model under most circumstances. Decision makers should work closely with subject matter experts to define acceptable accuracy levels for the AI models. Decision makers should also understand and document the inherent risks in trusting these AI models. They should make sure that all stakeholders know data is inherently biased and so are the models, and measures should be taken to mitigate these biases.

Pain Point	Description
Budget and Time to Market	Whether new AI solutions are customer focused or internal to the business, decision makers need to understand the time to market to build and rollout AI solutions. Considerations such as how long does it takes to build AI models, how to start a proof-of-concept (POC), prove its value, and operationalize need to be answered. The cost associated with such initiatives should be estimated. In addition, assessment should be made to see if solution providers and vendors can help. In general, decision makers should identify and estimate the costs associated with managing and operating data and AI models.
Personnel	Decision makers will have to identify the existing skillsets of their personnel and see if they require additional training. Decision makers will also need to determine if new job roles are required that can be filled by their personnel, potential new hires or outsource to IT providers to help businesses in their AI endeavors.
Regulations, Privacy and Ethics	It is paramount to adhere and conform to all applicable legislation and regulations related to consumer data privacy. Thus, decision makers will have to ensure that clear data ownership is defined, that the data is collected and managed in accordance with the laws and regulations, and that the data is secure. They should seek expertise and help establish general data standards and practices that can be easily audited. Decision makers should understand the tradeoffs when building AI models that require transparency.
Selecting the Right Implementation Partner	Partner selection will depend on current state of AI planning within the business. Looking at expertise across the technology, people and processes will enable decision makers to begin an engagement to support strategic objectives and provide clear focus. Deep data science experience, domain expertise, past performance and scale are important for long-term success.

In the following sections, we dig deeper into some of these pain points from a decision maker perspective and offer guidance and applicable best practices.

Avoid Shiny Object Syndrome and Focus on Business Use Cases

There is increased pressure on business executives and decision makers to incorporate AI into their product offerings due to the buzz and hype about AI in the market. Businesses that have already built AI-infused products promote and market their use of AI to differentiate themselves from their competitors. Executives and decision makers thus wanted to build AI-infused products yesterday. But there is an elevated risk for businesses to simply treat AI as a magic wand or a cure-all elixir without studying and analyzing—with due diligence—how AI can provide measurable value to their business and a good return on investment.

According to Gartner's 2021 Top Technology Trends Report, AI appears in multiple categories of strategic investment that decision makers need to consider—just like cloud computing was a decade ago. However, depending on the business requirements, available resources, and the type of data available, AI may or may not be the answer to a business's current problems. Simply leveraging off-the-shelf business analytics and intelligence solutions may address the immediate needs of a business without the need to invest in a complex and resource intensive AI project.

Self-Evaluation

It is common knowledge that AI involves large datasets, and multiple tools and techniques for it to provide meaningful value and make predictions about future business outcomes. Before you, as a decision maker, embark on the AI journey, it is recommended that you do a self-evaluation based on these questions:

- Do we have a clearly defined scope on how AI can solve our business challenges?
- Do we have the right level of executive sponsorship within our company to support the data and AI initiatives?
- Do we have the appetite or budget to make investments in data and AI initiatives?

Best Practices

If the answer is yes to the questions above, you can start identifying and prioritizing your key business use cases that will benefit from AI infusion. Here are some recommendations:

- Be extremely specific on the scope and desired outcome for your AI projects.
- Ensure that you have the buy-in from all key stakeholders (including your C-suite) to implement AI projects.
- Always start with a narrow scope and vision that may not even use AI and

incrementally build more complex AI solutions. Simple things sometimes bring the most value.

- Research case studies from industry leaders and vendors to clearly articulate your expected outcomes.
- Review industry white papers and use cases to validate your proposed AI projects.

Resources

Instead of always starting from nothing, review existing resources like case studies, blogs and thought leadership materials.

- Review analyst reports like Gartner, Forrester, IDC, etc., and materials on AI business cases.
- Review Gartner Magic Quadrant and Forrester Wave to identify the leaders and innovators offering AI enabled solutions. It can give you a sense of the strengths and weaknesses of these various solution providers and vendors. This can help you understand what to look for from a business case perspective.
- Each provider puts out case studies with their clients on success stories. Review them for inspiration and direction but use caution. While useful, there may be a bit of an oversell.

It's All About the Data and Modeling

As the pace of change continues to accelerate, there is a divergence between organizations that successfully leverage their data as a strategic asset and those that do not. A 2019 Wall Street Journal article highlighted that most corporate executives felt their AI efforts would be hindered by poor data quality. The top challenges to leveraging data across the enterprise, according to executives surveyed by PwC, include:

- Poor data reliability (34%)
- Inability to address new regulations affecting data protection and privacy (33%)
- Inability to adequately protect and secure data (32%)
- Data siloing or lack of sharing (31%)
- Lack of analytical talent (31%)
- Information systems not ready to exploit data (28%)

Leveraging data to make decisions may sound obvious, but many companies still have organizational and process silos that inhibit them from effectively leveraging their data.

Inconsistent processes and data silos extract a massive, sometimes hidden cost that personifies itself in suboptimal decision making and missed opportunities. These inefficiencies are compounded when organizations look to leverage AI/ML solutions since data is the fuel required to train and run the algorithms. Too often, data scientists and engineers spend considerable time and resources cleansing data before an AI/ML model can yield tangible benefits, and this can lead to elongated timelines, increased costs, and overall missed expectations.

83% of organizations see data as an integral part of their business strategy, yet 69% say inaccurate data continues to undermine their efforts.

—*Experian, Embracing the Data Challenge in a Digitalised World*

Best Practices

The following best practices have been identified to help enable better usage of data and unlock AI capabilities across your business.

- Establish alignment between corporate strategy and required data. It sounds obvious, but many organizations fail to explicitly align their corporate strategy with their technology and digital strategies, including the supporting data needs. Doing so helps you clarify the gaps that need to be addressed and relative priorities.
- Define the data boundaries and owners. Access to the right information at the right time is the overarching goal of efficient IT systems, and there needs to be a data strategy that supports the corporate strategy. There are several data artifacts and methodologies that help define the data ownership, usage, and other aspects around data management. Large organizations may have a centralized data or analytics group, but it is critical to map out data ownership by organizational groups. There are new roles and titles such as data steward that help organizations understand the governance and discipline required to enable a data-driven culture.
- Leverage and upgrade your existing data management solutions. Commercial software and packaged solutions are now offering inbuilt AI and advanced data capabilities. Leverage these data management and business intelligence platforms to see if they can fulfill your business needs.
- Ensure you have the right sets of tools for model development, deployment, and lifecycle maintenance.
- If applicable, reach beyond your current internal systems and data. One of the biggest drivers of innovation and efficiencies is enabled through external data sets. Explore how you can get data from your customers, how you can share data with your trading partners and vendors and look for ways to automate data sharing beyond your organizational boundaries. There are several data marketplaces where you buy data sets, and these new, external data sources can lead to significant insights or products to enhance an AI/ML model. Of course, see to it that you can trust these data sets.

Budgeting, Proof of Concept, and Rollout

As a decision maker, a key question to consider is the time it will take to rollout a functioning and operable AI-infused solution. You should think about what it takes to do a proof-of-concept (PoC), measure and prove its value, and figure out the budgetary requirements to build, deploy, and scale the solution to production.

AI projects typically take anywhere from 3 to 36 months depending on the scope and complexity of the use case you are building. There are certain open-source tools and libraries as well as machine learning automation software that can help in accelerating this cycle. Also, before you embark on the project, you will need to ensure that you have access to "clean data," i.e., data in a format that is easily accessible and fed into a machine learning platform. Often business users underestimate the time it takes to do "data prep" before a data science engineer or analyst can build an AI algorithm.

Once you have identified a project or a business challenge, you can begin planning for a PoC, which will include data sources, technology platforms, tools, and libraries to train the AI models leading to predictions and business outcomes. It will take multiple iterations for a data scientist to deliver 90%+ probability, which may be required to prove value for the model you are building.

Lastly, more than 80% of the AI projects typically do not scale beyond a PoC or lab environment. Customers often face challenges in standardizing their model building, training, deployment, and monitoring processes. You will need to leverage industry tools that can help operationalize and scale your AI pipeline. This is typically referred to as machine learning operations or MLOps.

Self-Evaluation

Here are a few questions to consider as you embark on your first AI project:

- Do I have all the data sources I need to start the PoC?
- Do I have all the required talent, tools and technology platform to start my PoC effort?
- Do I have checkpoints or stage gates in the process to evaluate success/failure of the PoC?
- What threshold or probability boundary should I set for the AI prediction model? For example, is 75% good enough or should it be at 90%+ accuracy?
- What timeline is my organization willing to give me to scale from PoC to production?
- What kind of expenses should I think about to implement the best practices? Expenses could be related to infrastructure, data acquisition, crowd/SME annotation, and so on.

Best Practices

- Review your projected business value from an AI project over a 12-to-36-month time horizon.
- Consider the total budget including internal headcount, contract resources, IT infrastructure (including application and cloud licenses) to calculate your total investment on the AI project.
- Cost of data acquisition (sourced internally or externally) should be reviewed as part of the total budget.
- Be realistic in setting your expectations.

Workforce Skills Upgrade and Hiring

Businesses should take stock of the current team and then consider a people strategy. The strategy should consider reusing and repurposing existing staff and upskilling, as well as training, hiring, and bringing in experts for temporary consultation. Some businesses might simply need their third-party IT service partner to provide the AI skills. And as the organization matures, several new roles can be considered with the help of a data-driven culture—data scientists, data engineers, and others. In addition, depending on the size of the organization and needs, there may be a need to build groups and communities to help the overarching AI infusion goals. For example, an AI community of practice and excellence or a cross-functional automation team.

Self-Evaluation

Here are a few questions to help you evaluate your personnel readiness:

- Do we have any data scientists on staff, or do we plan to hire in the near future?
- Can we leverage some of the developer expertise within the IT team to support AI development efforts?
- Do we have machine learning architects/data engineers who understand the data silos and can develop programming models?

Best Practices

- Hire data scientists and AI engineers if your budgets permit.
- Train full-stack engineers on AI and help them upgrade their skills.
- Procure corporate subscriptions to online training platforms and put together a tailored course curriculum.
- Host meetups and talks from senior practitioners from the industry to improve the exposure on best practices.
- Encourage engineers to go to industry and academic conferences to increase their exposure and understanding of the latest technologies, and platforms.
- Create centers of excellence for various areas such as data management, data science, AI/ML operations, and model evaluations.

Selecting the Right Implementation Partner

While implementation of your business's AI-infused solutions should be led by practitioners, debunking myths early on would help you, the decision maker, make these AI solutions successful. One of the common myths is that clubbing together a mix of open-source tools and deploying them in the cloud is a sure-shot recipe for a successful AI project. In reality, that is only partially true, and only if you are creating a small PoC for kickstarting your AI journey.

AI projects are complex; however, implementation does not have to be. There are leading industry vendors available today to assist you in your AI journey. But you need to be cautious in your vendor/implementation partner selection. There are hundreds of startups and emerging vendors that may not have the resources or investment capital to sustain in the long run. Bear in mind that AI solutions are not cookie cutter and while most solutions available today will meet 80% of your requirements, you will still need to work on customizing the remaining 20% to suit your business needs.

For starters, begin by researching use cases and white papers that are readily available in the public domain. These documents often mention the types of tools and platforms that can be used to implement business specific AI solutions. Explore your current internal IT vendor list to see if they have offerings for AI solutions within their portfolio. Often, it is easier to extend your footprint with an incumbent solution vendor than introducing a new vendor. Once you have a shortlist, feel free to invite these vendors to propose solutions to meet your business needs. Based on the feedback, you can begin evaluating and prioritizing your vendor list. Lastly, AI solutions are not a single widget or an application but an ecosystem of solutions that include infrastructure, data sources, tools and frameworks, libraries, visualization apps, etc. These can be consumed in the cloud (public or private) or within your existing datacenter or in a hybrid landscape.

Best Practices

Here is a checklist of questions to ask your implementation partner:

- How long have you been offering AI solutions?
- Do you have any use cases/examples you can share that aligns closely with my industry or organization needs?
- Do you have alliances with multiple vendors and an ecosystem to deliver me a complete solution?
- Do you have the resource bench to assist me in deploying this solution across the enterprise, and across geographies? Probe for data scientists, data engineers, machine learning architects, service engineers, etc.
- Can you integrate the AI solution within my existing IT footprint?

Part 2: Practitioners

The confluence of big data, the use of GPUs, and advances in algorithms have now made the field of AI mainstream. In some use cases, the accuracy of AI model prediction is either approaching or is better (in some cases) than humans. The mainstreaming of AI technology is also leading to its adoption for solving real business problems, such as customer churn prediction, X-ray diagnosis, fraud detection, autonomous cars, etc. In this section, the CompTIA AI Advisory Council defines the practitioner personas that must cooperate to develop and deploy AI solutions and describe the phases of end-to-end AI application development lifecycle. Each of the phases includes challenges being currently faced, the best practices that are being employed, and references to some useful collateral material.

AI Practitioner Personas

There are many AI practitioner personas that must collaborate to develop and deploy an end-to-end to AI solution lifecycle. However, as the industry continues to evolve, more specialization is taking place and new personas are emerging. Here are some common AI-related personas:

Data Scientist: Data scientists evaluate what data sources are required, define/perform data cleaning tasks, experiment with different AI algorithms and develop AI models. They determine the adequacy of an AI model in conjunction with subject matter experts. They usually rely on data engineers to set up the AI systems. Presently, many AI models and algorithms have been pre-packaged and are readily available. Thus, even non-statistics/computer science college majors are able to perform AI model-building operations by leveraging these prepackaged AI algorithm models or libraries. However, an advanced college degree could be required when trying to develop new AI algorithms. There is also a need to label raw data and make it ready for use by AI algorithms. The data-labeling tasks can be performed by humans, automated systems, or in some cases, by data scientists themselves. Due to the size of the datasets involved this is not a trivial task and requires a group of humans or automated tools.

AI Engineer: AI engineers are trained to construct or build the system. They worry about wrapping the algorithms as services and making them available via software-as-service (SaaS) or on-prem software that is packaged as part of the product installation or service access by implementing security, scalability, multi-threading, redundancy, and high-availability aspects of the solution. These are software personnel who have familiarity and dexterity with configuring and setting up AI frameworks. They typically install the right versions of different software and manage the AI process pipelines in conjunction with IT staff. They work closely with data scientists to ensure their computation needs are being properly met. Increasingly, AI engineers are also taking on the role of data scientists as AI/ML algorithms are increasingly available as part of open-source libraries.

Data Curator: AI models are only as good as the data that is being used to train them. As enterprises are building scalable systems, there is a need to procure data from multiple sources and ensure it meets compliance, privacy and quality standards. Historically, data scientists also performed this task in conjunction with IT personnel. However, now there is a need for a specialized persona to coordinate with legal, security, IT, finance, data scientists and DevOps teams to make the data available for data scientists. Currently, many large organizations have a team that is dedicated to quality assurance tasks related to data.

Here are some additional traditional personas who collaborate with the AI-specific personas:

Subject Matter Experts: Data scientists typically do not have deep expertise in a particular subject area. Similarly, subject matter experts do not usually have deep expertise with data science algorithms and tools. Thus, these two personas must collaborate to get a working AI model.

IT Staff: IT staff are involved in managing the data systems and AI systems that are used by data engineers and data scientists. Furthermore, they are playing key role with respect to streamlining security and operation processes as AI solutions span across public and private clouds.

DevOps Engineers: These software engineers are responsible for building the business application that integrates with the AI modules. They are full-stack developers who are responsible for programming the user interface, application server and backend systems. In many organizations, they have an additional application architect persona who collaborates with DevOps engineers to architect, develop and deploy software.

Program/Product Managers: This persona is responsible for the overall management of a product or a project within a company. They collect business requirements and work with the above personas to get a set of technical requirements and help coordinate the overall the execution of the project. Historically, they have experience in dealing with regular software development projects and only now are beginning to get acclimatized with the lifecycle of AI-infused projects.

AI Lifecycle from Practitioner Standpoint

The AI lifecycle consists of the following phases: requirements assessment, data processing, model building, model deployment, app integration, and monitoring and optimizing (see Figure 1). There are different people involved during these different phases. In this section, we describe each of the AI lifecycle phases, the challenges being encountered by practitioners, and some related best practices.

Figure 1: AI Lifecycle

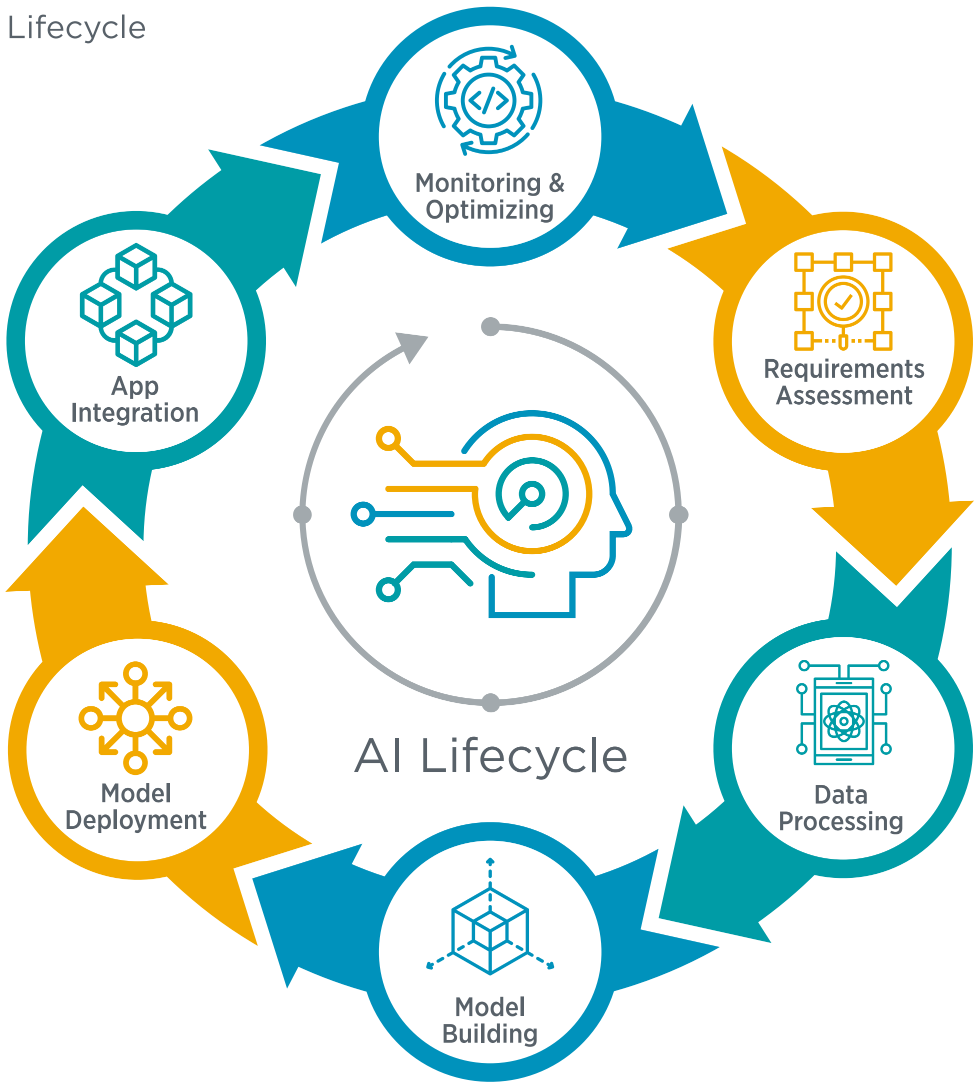

Phase 1: Business Needs Assessment

During this phase, an assessment is made with respect to the business objectives and technical requirements of an AI project. Typically, an AI project is part of a bigger application that is trying to optimize business processes, provide better customer service, or predict threats and opportunities for the business. Typically, in this phase, a combination of business leaders, program/product managers, data scientists and subject matter experts work together to finalize the business and technical requirements for the AI project.

Challenges

Presently, in many organizations, business leaders and product managers who lack AI experience are creating AI product roadmaps without a clear understanding on what AI can and cannot solve. Due to industry hype, they are under pressure to leverage AI. In many cases, they can get small proof of concepts (PoC) quickly built and deployed in clouds, but they are having difficulty building scalable, production-grade, private/hybrid AI solutions due to lack of talent and absence of a scalable data management strategy. The presence of a proper data management strategy is a key prerequisite for the success of a production-grade AI solution.

Best Practices

Organizations and business leaders need to involve data scientists from the very beginning as they plan how to leverage AI technology as part of their product or service roadmaps. Most AI models and solutions need to be vertical specific to be relevant, making it important to ensure that the external consultants and AI technologists being hired have experience in building scalable AI solutions and have vertical-specific experience with respect to the AI models, available data sources and algorithms.

Phase 2: Data Processing

This is the costliest and the most time-consuming phase in the AI lifecycle, encompassing 70% of overall AI development cost. During this phase, an organization needs to procure the appropriate datasets (both internal and external), clean and label the data, perform quality assurance on the data with respect to compliance, data quality and data currency. In this phase, practitioners are getting data ready to be used by AI algorithms to build the AI models. A combination of data scientists, data engineers, IT personnel, data/compliance officers, security officers, and infrastructure staff work together to configure and create the appropriate data management systems and pipelines.

Challenges

The accuracy of the AI models is only as good as the data that has been used to train and build the AI model. Lack of access to quality (well-labeled, compliant, clean) data is one of the key challenges that most enterprises currently face. Furthermore, in many cases, they need a scalable data management platform to manage the data that is being used for AI model training. Privacy, access control, and compliance with respect to government regulations, such as GDPR and the California Consumer Privacy Act, are some additional challenges that enterprises face with respect to how they handle data as part of their AI solutions. Because 75% of enterprise analytics applications on average use around 10 external data sources, enterprises need a well-defined governance model for procuring and handling these external data. Currently, 74% of business users still use spreadsheets to exchange and manage data, while 69% of users rely on email attachments to exchange and manage data, according to Ataccama's Enterprise Data Readiness Report.

Best Practices

In addition to hiring data scientists, organizations need to hire architects who can design scalable data lakes and warehouses. Organizations need to also have dedicated groups that procure external data legally and ethically in a cost-effective and timely manner, perform data labeling and cleaning, and do quality assurance on data quality. Organizations need to also move from ad hoc data management by data scientists to production-grade repeatable processes for handling data management tasks. Many data governance and management tools for doing data management tasks are now available in the market.

Phase 3: Model Building

In this phase, an AI model is built by leveraging various types of algorithms and data. An AI model is usually a digital representation of an artifact (i.e., an airplane, a car, a factory, etc.). Typically, by using these models, practitioners analyze what happened in the past (i.e., failure analysis), assess a present situation in real time (i.e., video surveillance), or predict future state based on past/current trends (i.e., a stock price). AI algorithms can be:

- Traditional, rule-based systems where experts specify rules about how the system should behave when it encounters different events.
- Machine learning systems that derive patterns from large datasets (past state) based on different statistical techniques.
- Deep learning systems that leverage large datasets, computation power, and different types of neural network algorithms for prediction.

Data scientists experiment and iterate with different types of algorithms and datasets in order to come up with an accurate model. Model building is very resource intensive (compute, storage, networking) and in many cases, people use models built by other companies as a starting point instead of building a model from scratch (i.e., leveraging an image detection model that has been built by a public cloud after analyzing millions of images). Data scientists work with data engineers to get model training systems configured, and they work with subject matter experts to validate the trained models.

Challenges

One of the main challenges with the AI model training/building phase is that the data centers of many organizations cannot easily host the AI training hardware due to the high-power density (30KW for a fully loaded rack) requirements. Second, in many cases, data scientists are not able to explain how an AI model works in a way that is easily understood by non-technical stakeholders. This is especially true for deep learning algorithms. Another key challenge is how to ensure that the data being used to train or build an AI model is not biased because biased AI models can lead to ethical issues and impact sound business decision making. Finally, in some use cases, organizations are moving massive datasets to a central location from edge locations for creating their AI models. This mode of operation is quite expensive and time consuming.

Best Practices

Many organizations are doing model training operations either in public clouds when they are willing to move their data outside their security perimeter, or in their own premises in a co-location facility because these locations can satisfy the high power requirements of specialized AI training hardware (i.e., GPU-based hardware). Secondly, many organizations are leveraging and customizing pre-built AI models that have been built by public clouds (i.e, image detection, audio to text translation, etc.) by shipping their private data to the public clouds for model training. Thus, they are leveraging the AI innovation and data sets that are present in the clouds. They subsequently ship the trained model back to their data private data centers for inferencing. This hybrid model allows organizations to leverage the best of public cloud innovation while keeping the execution of their overall application private for privacy and compliance reasons.

Phase 4: Model Deployment

Once a model is built, it gets deployed and used for inferencing. The term "model inference" is used to refer to when the built model is being used for predictions. Model inference phase is not as resource intensive as the model training phase. Field programmable gate arrays (FPGAs), GPUs, and CPUs are the different types of hardware that are used for model inferencing. Typically, data engineers along with IT personnel deploy a model as part of production workflows. There are AIOps and AI PaaS frameworks that help to manage the model training and deployment process.

Challenges

Historically, model training and model deployment have been performed at the same location (i.e., in a central location), but increasingly for latency, cost and privacy reasons model inference is being performed in a distributed manner at the edge close to where the data is getting generated. This, in turn, presents new challenges to an organization because now they must design distributed AI workflows that span across multiple geographic locations (potentially spanning across multiple security domains).

Best Practices

Organizations must adopt distributed AI architectures where model training and inference can occur at different locations in order to deal with large data sets that are being generated at the edge. As an industry, we are now entering the era of federated analytics where model training and inference will occur in a distributed manner at different locations for privacy, performance and cost reasons. Enterprises are beginning to leverage distributed/federated AI PaaS frameworks for deployed distributed/ federated AI applications.

Phase 5: Application Integration

AI model usage and functionality is usually part of a larger application that is trying to solve a business problem for an enterprise. Data scientists and engineers need to work with subject matter experts, IT departments, and traditional software architects and DevOps personnel.

Challenges

The software development and quality assurance process for AI development is different than for traditional software development because of non-determinism in the AI model training and inference process. That is, as the data sets change (sometimes even different runs with the same datasets), the AI training and inference process can yield different results. Furthermore, in many cases, developers cannot explain why the model has predicted a certain outcome, and the quality assurance people have to design their testing processes to test for goodness of the predictions rather than whether the software works or not (binary mode). Finally, many organizations want to leverage the innovation and pre-built AI models in the public clouds while they also want to keep the persistent copy of their data or run the non-AI portion of their application in a private cloud. This leads to hybrid AI architectures that are complex to manage with respect to security, performance and management.

Best Practices

Typically, for a small proof of concept, a data scientist can leverage the AI frameworks that are available in public clouds to showcase a working prototype. However, for building and integrating AI services as part of a larger business application, it is necessary to either train existing software professionals or hire AI application development experts or consultants. In many cases, enterprises should also integrate data scientists as part of regular software DevOps teams rather than having them siloed as part of a separate data science team.

Phase 6: Monitoring and Optimizing

After an AI model is deployed and being used in production, it must still be monitored for accuracy because often the data that was used to train the model has either changed, there access to better or newer data sources, more data of the same type is available, or the application requirements have changed. Therefore, the AI lifecycle process described is cyclical. In some cases, the AI workflow is completely automated, and in other instances, product managers have to sign-off if major changes are being made to the underlying data or algorithms.

Challenges

Similar to how automation tools and frameworks have automated the software development process in order to accelerate bug fixes or bring new features at a faster cadence, there is a need for AI operations frameworks that automate end-to-end AI workflows. There is a need to automate data collection, data cleansing/labeling, model training, and deployment processes. Absence of this automation will make AI solutions stale or brittle and will lead to AI model prediction errors.

Best Practices

A lot of work has recently been done in the area of AI operations pipeline management. Frameworks are now helping enterprises define and automate their AI pipelines and make them production-grade ready. With the advent of cloud native technology era (i.e., containers, Kubernetes, etc.), it has now become easier to put various AI software components (i.e., data cleaning, data management software packages, AI training platforms, AI models) into containers and continuously upgrade and deploy them across distributed locations.

Conclusion

AI technology is now being leveraged as part of business-critical applications. The industry has moved on from demonstrating toy AI solutions. It is now ready for deploying robust business solutions that leverage AI technology. As part of this transformation, infusing AI into a solution requires the participation of business decisions makers and practitioners from the beginning of a project. In addition, more specialization is taking place with respect to the personas who are helping build these production-grade AI solutions. There is also a proliferation of AI tools and services to help enterprises manage their data and build scalable AI solutions. Decision makers and practitioners alike can use the best practices articulated in this guide to build successful business applications that leverage AI technology.

Resources

Artificial Intelligence (AI) Terminology: A Glossary for Beginners

Business Considerations Before Implementing AI

Emerging Business Opportunities in AI

How to Leverage AI to Get the Most from Your Company's Data

About the CompTIA AI Advisory Council

The CompTIA AI Advisory Council brings together thought leaders and innovators to identify business opportunities and develop innovative content to accelerate adoption of artificial intelligence and machine learning technologies. The council is committed to building the strategies and resources necessary to help companies leverage AI to be more successful and collaborates with CompTIA's other industry advisory councils to further study the IT channel, blockchain, drones, business applications and internet of things.

CompTIA.org

CompTIA is responsible for all content and analysis. Any questions regarding the report should be directed to CompTIA Research and Market Intelligence staff at research@comptia.org.