issue 2

AI Adoption to Revolutionize the Future of Business

Implementing strategic & scalable AI and ML projects

Gartner

Executive Summary

Artificial Intelligence along with Data & Analytics has been the most disruptive technological advancements businesses are experiencing. The impact that AI has made with its immense computational power to analyze large volume and variety of data at great velocity has allowed businesses to draw meaningful insights from their data sets for business gains. From the automation of uneconomical applications to create new platforms, products, and offerings, organizations are using AI to harness data deliver extravagant consumer experiences.

While we have seen unprecedented advances in AI and Data & Analytics space, it still presents businesses with unforeseen difficulties and challenges repeatedly. As many as 60 – 70% of Big Data and AI projects fail to take off due reasons ranging from poor planning to lack of resources. In some cases, the scalability takes the center stage in the absence of practical test cases and lack of data. These are just some of the challenges that the data and analytics leaders of modern day are faced with. While some are navigating their businesses towards successes, some are learning their way out of these challenges.

We have faced our own fair share of challenges & difficulties in devising successful, meaningful, and scalable AI projects. Our experience in the AI and Data & Analytics space has allowed us to create the principles & guidelines that can help any industry with huge data sets to create & scale AI projects, with the desired ROI. In this paper, we look at some of the hurdles that obstruct the scalability of AI and ML projects. We will also be discussing some architectural best practices, which will help businesses build elasticity in an otherwise fragmented AI landscape.

Abstract

There is no business area which has been potentially untouched by AI. The ever widening scope of AI is one of the largest breakthrough in Science. While AI and Machine Learning are being widely adopted across businesses for not just insight generation but for automation and embedded intelligence too, there is still a whopping set of statistics that indicate that they may not always deliver the extravagant promises originally envisaged. As many as 60-70% of Big Data and AI projects fail to take off beyond preliminary experimental stages.

As AI tends to become everyone's business, pervasiveness of all forms of machine learning also leads to local pockets of intelligence in an enterprise. While the individual use cases generate a lot of internal data monetization, the value is not entirely realized due to lack of scale. Often times, the use cases reside in siloed data lakes and are rarely adopted main stream to become embedded within business processes.

While there are several reasons to failure, there are also a whole set of AI projects which exponentially succeed and scale to newer and more enhanced areas. There are a certain set of guiding principles that help scale AI projects and lead to true realization of ROI. In this paper we will explore potential reasons which impede the scalability of AI and ML projects. We will also touch upon some architectural best practices that can help build elasticity in a usually fragmented AI landscape.

Source: Tech Mahindra

Source: Tech Mahindra

AI and ML in business

Taking a step back let's look at some of the predominant use cases across different businesses where AI is thriving and revolutionizing the marketplace. Retail and E-commerce are completely transformed through AI powered solutions. Providing product suggestions, leveraging chatbots and virtual assistants to help customers find what they need, radical personalization to drive better customer experiences are just some areas where AI is widely utilized today.

In Manufacturing many of our customers both in auto and aerospace are using IoT analytics to drive preventive maintenance. One of the leading aircraft manufacturers is using AI to predict remaining brake life of aircrafts and hence better address maintenance needs. Cognitive Repair Solutions are using NLP based searches to qualify the repair and identify a solution. We are also implementing ML approaches to optimize the engineering parameters to arrive at an optimum design solution leading to shorter time to market.

AI for drug discovery and development is gaining tremendous foothold and it may not be too far beyond when Robot Scientists come into being. At TechM, for instance, we have RoboVigilance which can largely automate the process of pharmacovigilance. Medical imaging is a growing area where AI is being used to automate cell counting, image analytics, tumour detection etc. Robotic surgery is a prospective area which is being extensively researched and utilized. AI has a deep relevance to the area of personalized medicine and reinforcement learning is likely to be used with profound impact in this area.

Cybersecurity is increasingly being AI powered and is becoming more sophisticated and powerful. Detection of intrusion is being accurately predicted and instantly identified. Synergy of regulatory requirements such as GDPR and analytics are fast gaining grounds. While the beneficiaries of AI are very many to quote, there is a fundamental gap which exists in the mainstream adoption of AI across enterprise processes. The gap emanates primarily from the reliance on data till such time as intelligent training data can itself be generated and from the slowest moving part of a data landscape pertaining to data integration and management.

Why AI is not scalable today?

AI, we know for most practical use cases, relies heavily on availability of good quality, well-engineered data. A lot of that data today also lies outside an enterprise. It has become extremely critical to leverage orthogonal datasets in addition to standard enterprise data. In such a scenario we routinely encounter the following key challenges:

1. Development Cycles:

Orthogonal data can be in the form of externally purchased data such as market information, risk data etc. or from the social media to name a few sources. Coupled with complexity of extending enterprise data warehouses, it is an uphill task to integrate new data sources. For many of the enterprises there is still a heavy reliance of generating insights from the data warehouses. In such scenarios it is time consuming to adapt to newer datasets which inhibit the ability to explore new avenues through AI. The key to solving this issue is not just limited to a technology solution such as Big Data but also requires careful thinking around design patterns which help AI become extensible.

2. Data Acquisition:

Data being the foundation block of AI implementations, has to be able to augment an AI system adequately and hence until we have that level of data quality, we may not be able to produce the right results. Typically, enterprises have different sources for different needs and many of the times the granularity of data across these sources is too hard to homogenize. To further complicate, external data vendors may continuously update their feeds to support additional information. On most occasions enterprises may not have a well-defined data strategy for data quality and governance. Due to increasing federation of data, it is common to have several siloed data sources that are monetized locally or at best at a department level. Data Discovery prior to acquisition thus becomes a huge challenge. Even if there are big data lakes to expedite data provisioning, lack of discoverability, a firm Business Glossary and enterprise metadata leads to ambiguity during data acquisition.

3. Deployment Cycles:

AI systems can many a times be computationally expensive. While they may perform very well in limited settings, scalability on a larger enterprise canvas may be prohibitive. For instance, one of our customers who tried to implement forecasting at an SKU level succeeded immensely when the algorithms were run on a select set of products but when the prototype was extended to all product lines on a cloud infrastructure, the resulting AI system was computationally too expensive, time consuming and no longer viable. This is where the issue of scalability and elasticity comes in.

Another aspect of deploy-ability stems from the fact that when AI systems are stretched to wider datasets they may not adhere to the same level of accuracy. As a result, they need to be tuned and refined iteratively across multiple learning cycles before generalizing to a big picture use case. This process is often time consuming and needs a much longer deployment cycle time than originally estimated.

4. Evaluation cycles:

The statistical models in an AI system are trained (in most cases) on a given set of training data and one of the fundamental assumptions is that the behavior of the training data closely mimics the data on which the model will eventually be run. While enough care is taken to eliminate bias through a more iterative approach as cited before, it is fair to say that the resultant evaluation never has visibility to all possible outcome in a typical Unit Test Case scenario. That is to say that there are too many outcomes to test individually and hence evaluation itself has to be an iterative process. The models thus need to be tested and tuned every few years if not months to ensure ongoing relevance to the business.

Faculties of transformation

The centrepiece of transformations require us to ask appropriate fundamental questions to shape the strategic vision. What will data and analytics be used for? How will the insights drive value and which data sets are most useful for the insights needed? Many of the newer entrants into the data modernization journey often struggle to switch from legacy data landscapes to newer, nimbler and flexible architectures that can help in getting the most out of big data and analytics. Data modernization also needs digitisation of underlying operations and business processes more holistically in order to capture data from customer interactions, supply chains, equipment, and internal processes..

From a People perspective, the key to effective realization of ROI is also about finding the right champion for the AI use cases. The implementation calls for significant change management requirements and often means that people need to move away from their complacence with traditional BI systems. It calls for federating data and system ownership to end users and hence is often challenged by those who have organically grown the data landscape over a period of many years and sometimes decades.

From a technology perspective, there are too many options to choose from and evaluation of the appropriate platform including decisions such as build vs buy are critical to success of a data transformation. It’s also important to carefully consider the decision of implementing cloud based architectures as compared to on-prem ones. Sometimes, AI implementations are done on cloud for a variety of conducive factors such as elastic compute, on demand and low capex spends but these may soon lead to vicious cycles if the long term costs and architectural best practices are not followed. While cloud based AI implementations may be perfect in a fewer scenarios, they may not fit all.

How TechM is dealing with these challenges?

TechM has been working on a number of AI undertakings across verticals. Some of the examples are presented below:

- Market Mix Optimization - For both Life Sciences and other verticals TechM has implemented deterministic optimization models to identify appropriate marketing mix. We have implemented Mixed Integer Linear Programming to replicate the success across multiple customer scenarios
- we work heavily on IoT Analytics to identify preventive maintenance needs both in manufacturing and medical device scenario
- We have used AI in the context of Customer 360 to mine social media datasets, data generated from different communication channels, speech etc. We have helped enhance revenues through analytics driven lead identification, cross sell and up sell

- For retail customers we have implemented multivariate statistics models for improving revenue per customer by identifying the right targets and drivers for repeat visits in store
- We have also worked on massively parallel processed scenarios to truly deliver AI and ML at scale. Some of these implementations have enabled cloud based architectures supporting workloads of 16000+ machine learning model execution per minute

Appreciating the monstrous appetite of industry in the field of Data science (DS)- techM has built, PRISM - an IP, to help cut down DS life cycle tremendously. As much as robust data science platforms like PRISM important for enterprises, cataloguing applications assets which germinate across organization is the most profound activity to ensure the best tested AI/ML applications are made available to the needy. As we spoke in the beginning of this artefact, feeding relevant and curated data is another major step towards maximising ROI across an enterprise. This is where a metadata catalogue, accessible over intranet, crucially paramount for growing pursuits of data science professionals. TechM has built a home-grown IP, namely INFOWISE™, which comes to the rescue should there is a widespread need for enterprise metadata 360.

To address the needs of data quality menace, techM has blueprinted an approach called UNIFIED DATA MANAGEMENT FRAMEWORK, which guarantees prevention of “raw data” infections into enterprise applications and assets.

Review of traditional AI architecture style

AI Architectures need to be evaluated for the following key parameters to ensure long term success and their applicability to the business requirements:

- **Repeatability** – One off use cases may have a very narrow window of driving revenue potential or other forms of tangible outcomes. So when incepting an AI project, it is important to plan the generalized architectural ecosystem which can host a sequence of such use cases
- **Reusability** – It is important to consider reuse within and across departments or geographies. So building flexible and reusable frameworks is critical to achieving speed in AI implementations
- **Modularization** – Decoupling closely knit analytical capabilities and ensuring segregation of outcomes, as much as possible, will help in better integration of the AI system and ongoing adaptation to newer areas of an enterprise

Traditionally AI use cases have been implemented in a more monolithic manner where each use case was handcrafted to completely suit the purpose for and the data on which it was created. Typically, such implementations relied on discrete compartments of Feature Engineering, Model Engineering, Execution and Deployment. Each of these steps would be sequential and would have a sub system of its own. While they perform individual operations, such an architecture makes it difficult to cross leverage work and also make it harder to iterate across different steps of today’s AI implementations. Some of the drawbacks of this architectural style are:

- Discrete development silos each of which had heavy dependence on the preceding moving components
- Difficult to bring in new enhancements leading to longer development cycles
- Lower cross functional reusability leading to repetitive development and testing efforts

figure 1. Traditional reference architecture for AI

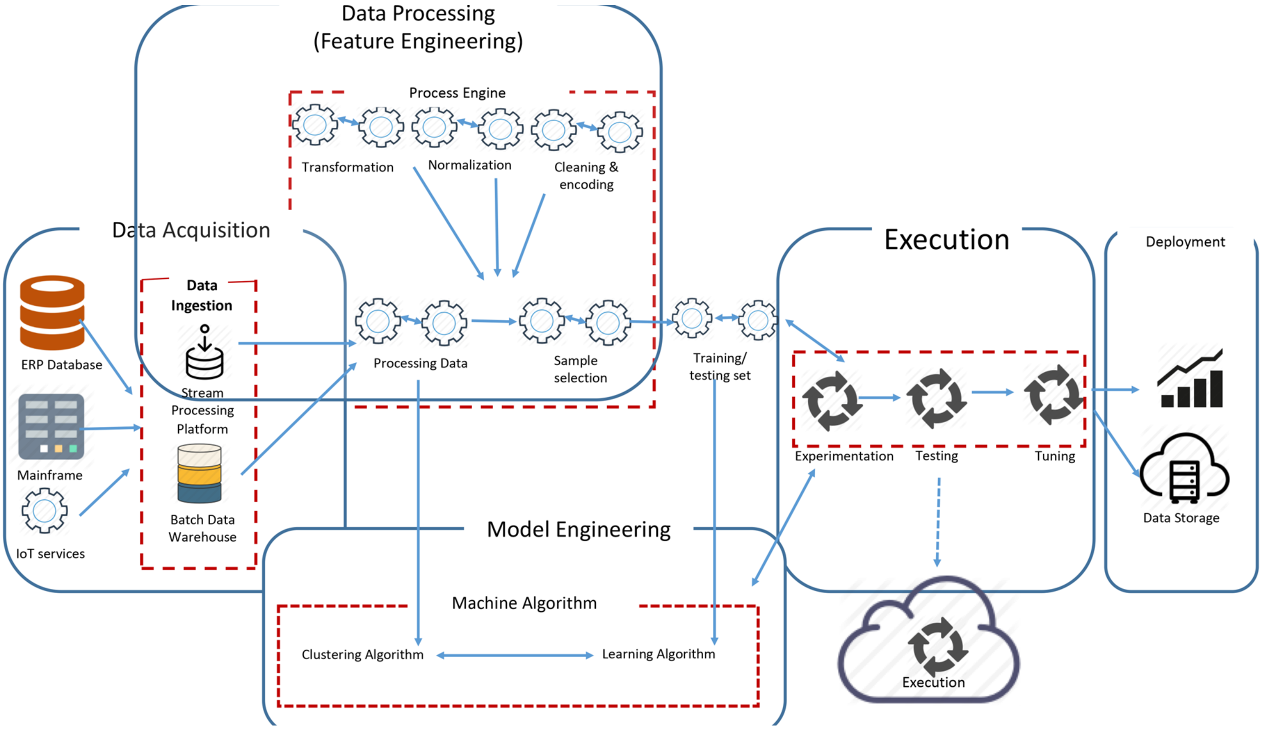

Source: Tech Mahindra

To mitigate some of the issues of this architecture, what is required is a system which is more easily extensible and supports Agile implementations of AI. The reference architecture should be more full proof to future scalability and enhancement needs.

Reference Architecture for AI at scale

Getting AI to work in a distributed architecture must be considered as part of scalable AI. Consider a situation where ML algorithms have to access data from dozens or even hundreds of servers. This would require significant scalability, elasticity and versatility. A reference architecture built for scale also needs to meet computaional requirements of advanced AI subjects such as deep learning. Engineers and stakeholders involved with deep learning may achieve better results from going deeper into data sets and manipulating them in more `overwhelming ways which would certainly need higher compute especially during development times. Deep learning projects are one example of how important it is to consider adoption of AI at scale. As more evolution continues to happen in this space, there will be higher pressure on machine learning systems to scale efficiently and elastically with time.

One of the reference architecture styles that can support such scalable applications is summarized below:

Apart from the standard data engineering and model engineering components, one of the key aspects of this architecture is its support to a wide variety of data and insight access components. The access layer could potentially support a Micro-Services driven architecture where each of the use cases could be run within containers and exposed as an analytical package for easier use and deployment.

Such flexibility could also help improve deploy-ability of the use cases on a variety of infrastructure both for on-prem, cloud, big data and non-big data architectures. The key best practices embodied by this architecture are:

figure 2. Reference architecture for AI at scale

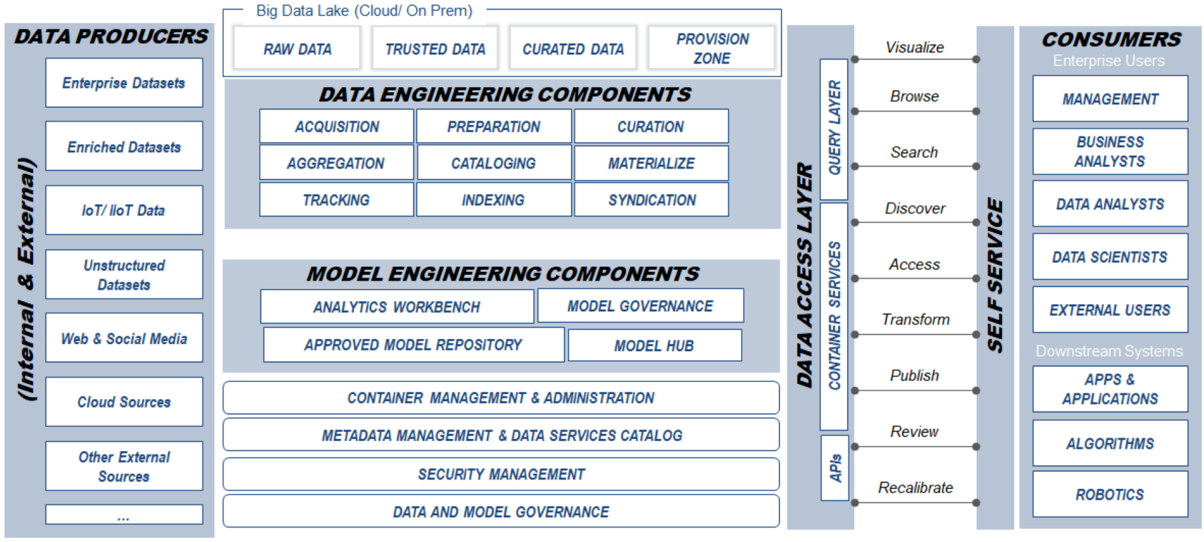

Source: Tech Mahindra

- Deploying workloads on an infrastructure that is purpose-built to handle the unique and growing demands of AI while providing the fastest time to training
- Large compute power to train models in less time
- High-performance storage to handle large datasets
- Seamlessly and independently scale compute and storage
- Handle varying types of data traffic
- Optimized costs

AI projects often experience bottlenecks during the training phase, when high I/O bandwidth with massive parallelism is required. Large data sets are important for increasing accuracy. AI deployments may start out small but may soon need to scale out to several petabytes. Additionally, performance needs may vary based on the training model used and the target application, requiring independent scaling of compute and storage. Designing a robust architecture enables independent scaling.

While the above section outlines the architectural and design best practices, bringing AI to scale also means process and people changes. From a process perspective some of the best practices include:

- Incepting a centralized Centre of Excellence for AI which will help in bringing much needed governance to an otherwise loosely coupled AI landscape. The central CoE can be made responsible for overall alignment of AI strategy
- AI projects are sometimes perceived as very short haul engagements but most use cases tend to evolve and sometimes may need a multi – year approach to replicating across multiple parts of the organization. It is important to set expectations and also identify interim milestones to ensure continuous business engagement
- AI projects must have a business champion since adoption may be a huge challenge otherwise. The business value and milestone wise ROI plan must be aligned and there has to be constant stewardship from the business side.
- AI projects sometimes are embraced in a more academic fashion and tend to be executed in R&D mode. This may inhibit the ability to realize full potential from AI projects. So it is important to also consider the overall enterprise strategy and make sure that mutual synergies are leveraged.

Source: Tech Mahindra

Source: Tech Mahindra

Conclusion

As organizations continue to mature on the data and analytics journey, it is critical to design systems, processes and change behaviours to suit the implementation of AI at scale.

The scalability issues also bring with them concerns related to both ethics, ease of handling data integration and security to name a few. Dealing with these concerns needs strategic thinking while use cases are conceptualized and must not be made as an aftermath of the implementation. From a technology perspective, the architectural styles need to be adaptive to changing needs of insight generation.

“AI pervasive across industry now-a-days and its potency is directly proportional to the depth of historical training data which it is augmented with. Human pursuit of excellence supplemented by AI takes businesses to leap frog into orbits which bring the aspects of impossibilities into the realm of reality. Thus, the confluence of AI with data & analytics not only optimize enterprise businesses but also generate innovative products, services and business models for the future. While almost all industry verticals have opened their doors to these disruptions, industries like manufacturing, Pharmacy, life sciences and banking are in the cusp of embracing those digital new-normal. We undoubtedly getting into an exciting next 5 – 10 years of digital renaissance.”

- KALYAN KUPPUSWAMY

Source: Gartner Research, G00355907, Whit Andrews, Tracy Tsai, Arup Roy, Balaji Abbabatulla, 3 July 2018

Overcome Obstacles to Your AI Projects

AI projects face unique obstacles, due to their scope and popularity, misperceptions about their value, the nature of the data they touch, and cultural concerns. To surmount these hurdles, CIOs should set realistic expectations, identify suitable use cases and create new organizational structures.

Key Challenges

- Artificial intelligence (AI) attracts intense interest from business leaders, due to media hype and the genuine value of automation improvements. This leads them to pursue projects with minimal input from the IT department and to have overly optimistic expectations for swift success.
- IT specialists and other workers express concern about the cultural, security, privacy and risk impacts of AI, while being skeptical about the benefits it will deliver.
- AI's disruptive profile polarizes leaders in IT and business units, creating opposing viewpoints that undervalue or overvalue its impact.
- Over 4,000 vendors of software and services say they are AI vendors.

Recommendations

CIOs seeking to develop initial AI projects should:

- Set expectations for a multiyear arc when pursuing AI, as projects generally extend for years before reaching the production stage. Benchmark, document and share interim achievements.
- Embrace a complete strategy, including formation of an AI center of excellence (COE), to mitigate risk, in collaboration with data and security, data management ethics and security officers. Involve these officers and executive sponsors in documenting and socializing realistic evaluations of risks peculiar to AI.
- Moderate inflated expectations for value and expediency by establishing an AI COE associated with, or as an aspect of, data and analytics centers. Include customer experience officers, security managers and executives.
- Require the AI COE to contribute to software and service vendor selection, training programs and hiring strategy, to make priority projects possible.

Strategic Planning Assumptions

By 2019, application functions based on artificial intelligence will be pervasive in 90% of enterprises globally.

By 2020, 50% of organizations will lack sufficient artificial intelligence and data literacy skills to achieve business value.

Introduction

Organizations seeking to power applications with AI face unique hurdles, such as the technology's complexity and immaturity and the perception that AI will have significant effects on people. AI is the latest (but surely not the last) family of technologies to be held up as a potential solution to any of dozens of use cases that have so far resisted technological resolution. That unreasonable perspective is itself an inhibitor to identifying the use cases for which AI is best suited.

A recent Gartner Enterprise Survey asked representatives of organizations that were at least planning AI projects to identify the three most significant concerns obstructing their progress. We sorted their answers into general categories (see Figure 1). Fear of not understanding AI's implications struck four in five organizations as one of their top three concerns. Finding a starting point was a concern for about two in three. Vendor concerns were not as immediate. Nor was enterprise maturity. (Because we asked for only the top three concerns, organizations may have had concerns they did not indicate, especially if those concerns were considered likely to emerge later in the process.)

figure 1. What Holds Organizations Back When It Comes to AI?

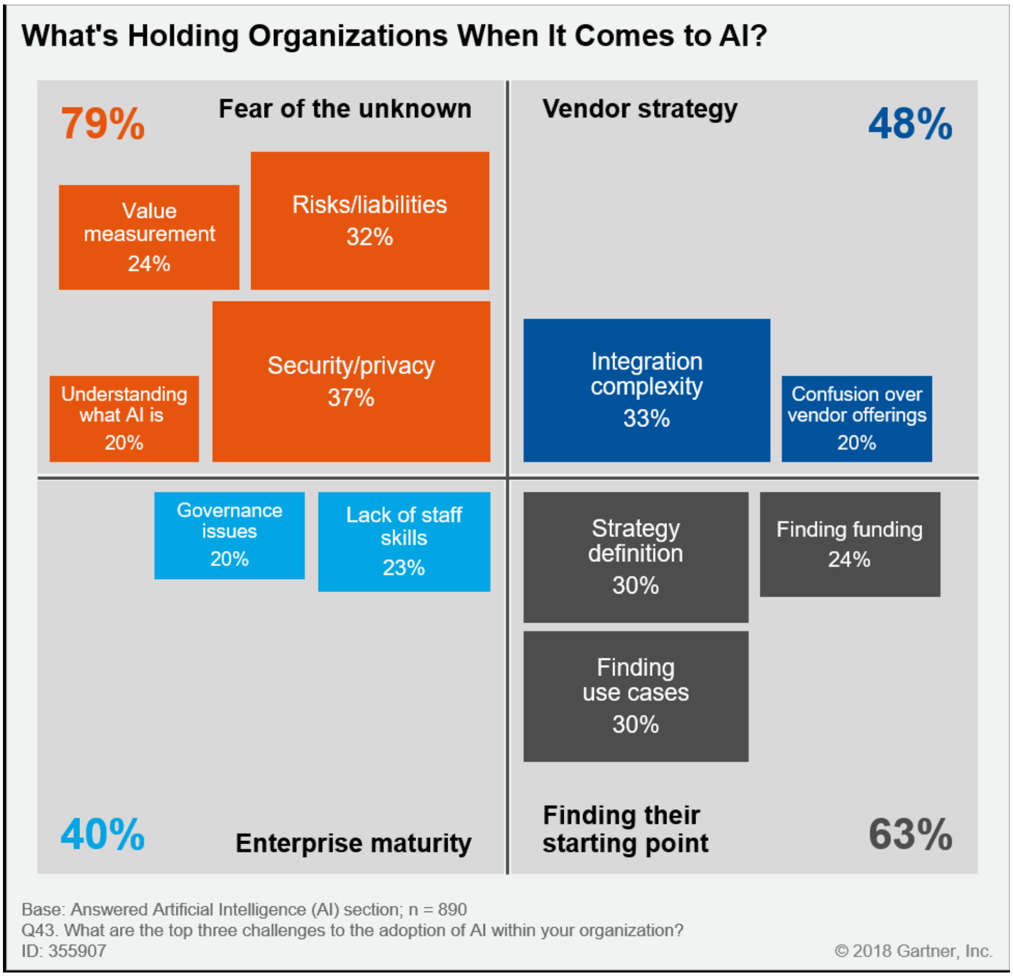

Source: Gartner (July 2018)

These obstacles are proving more challenging than organizations expect. While most organizations start out expecting to launch AI applications within two years of their first steps, in reality it often takes much longer (see Figure 2). Organizations that have passed the planning stage indicate, on average, that it will take them four years to achieve their AI ambitions with the application launch they intended. There is significant variation in project lengths. A handful of the organizations said it would take them less than six months, for example, and a similar handful reported it would take longer than 10 years. A larger number opted for between one year and five years. Such variation is normal and to be expected when considering new and challenging technologies. We expect the gap between expectations and reality to narrow with time as supporting technologies and products will emerge. Nevertheless, we recommend that organizations dampen expectations of swift impacts from AI applications, and tie them to the complexity and risk of the primary project they select.

To address these challenges, organizations must assign the right staff to projects, and concentrate resources by founding an AI COE. Alternatively, they may use an existing center for a related technology (such as data and analytics) with sufficient resources to support AI as well.

figure 2. Expected Duration of AI Projects Versus Reality

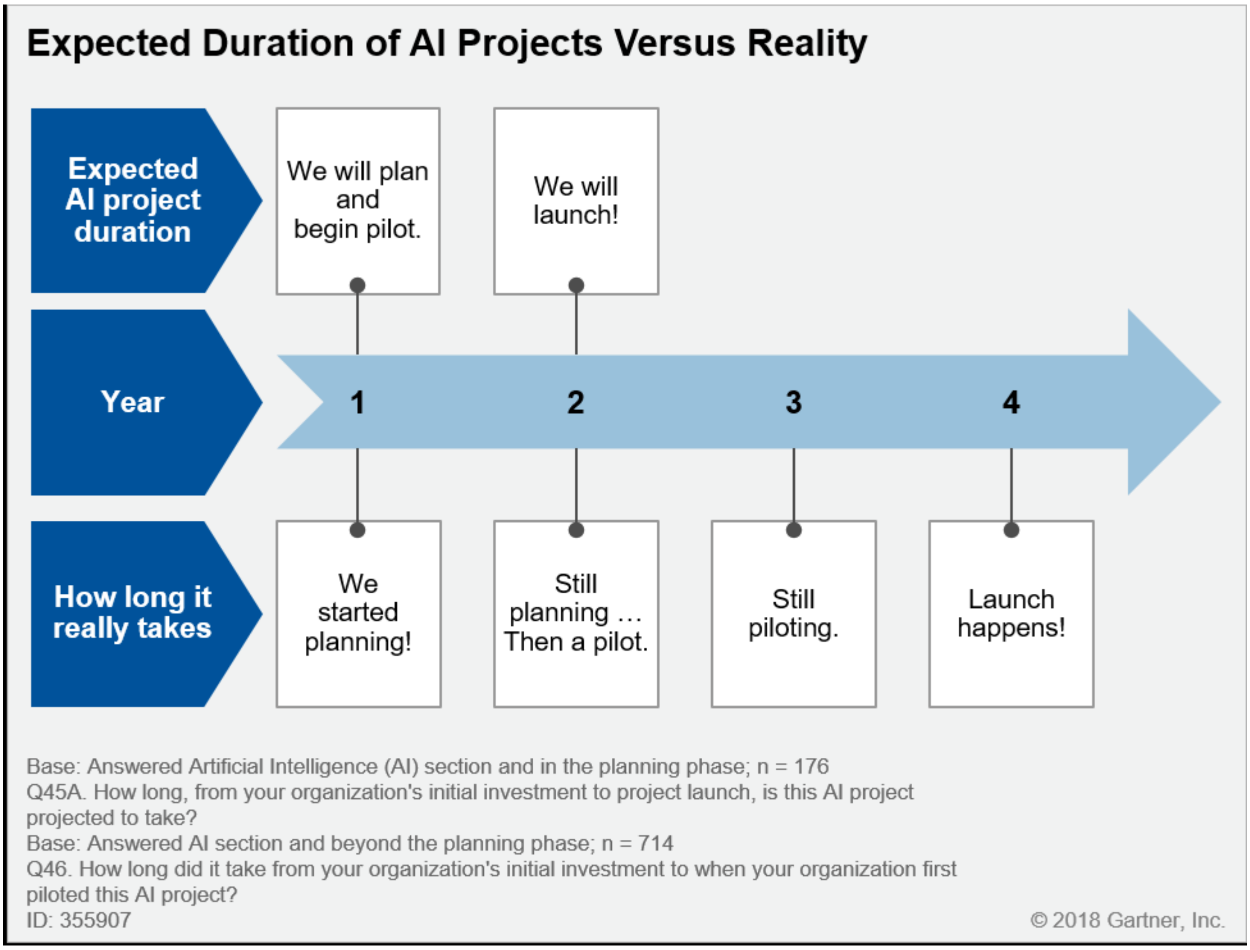

Source: Gartner (July 2018)

An AI COE should use the guidance of the chief data officer (CDO) and possibly include:

- Domain expertise
- Data scientists
- Business leaders (especially customer experience leaders), as well as human resources leaders.

Projects should have narrow aims initially and pursue thoughtful data preparation practices.

Analysis

Gartner's recent Enterprise Survey asked organizations what are (or had been) obstacles to their adoption of AI. Their concerns fell into multiple broad categories. Not all concerns were present in equal proportions. Organizations can mitigate most concerns by employing the right tactics (see Figure 3).

To overcome unreasonable concerns about AI projects:

- Centralize responsibility for AI strategy in an entity of limited life span, such as a COE for AI, or at least with roles within a competency center for data and analytics.

figure 3. Obstacles to Adoption of Artificial Intelligence and How to Address Them

Obstacles to Adoption of AI and How to Address Them

	Fear of the Unknown	**Finding Starting Point**
Phase of Impact: Starting	▪ **Impact:** Executives decline proposals; strategic paralysis; fears are various, from privacy to employment impact to data breach. ▪ **Tactic to overcome:** Emphasize partial automation and augmentation, not autopilot; demonstrate information management and data-handling maturity. ▪ **Intensity/relevance/duration:** Individual executive concerns can slow project momentum; outsource for skills and experience or target AI-powered application vendors. ▪ **Regional variation:** Most significant in markets with strong privacy concerns, such as Europe, the U.S. and China. ▪ **Industry variation:** Most significant in mature consumer-facing, labor-intensive, heavily regulated industries.	▪ **Impact:** Conflicting internal projects and delayed familiarity with AI potential. ▪ **Tactic to overcome:** COE can tally proposals and prioritize, set expectations and develop pilots; COE can justify funding. ▪ **Intensity/relevance/duration:** This obstacle tends to be subverted and dodged swiftly; the challenge is project governance. ▪ **Regional variation:** All regions face this challenge, with rough equivalence. ▪ **Industry variation: "**Digital first" industries typically start fluidly, as do the retail, higher education and financial services sectors.
	Enterprise Maturity	**Vendor Strategy**
Phase of Impact: Intermediate	▪ **Impact:** Staffing and governance plans hinder pilots and experiments. ▪ **Tactic to overcome:** Work with vendors, service providers and partners to improve readiness and skill availability. ▪ **Intensity/relevance/duration:** Because organizations must develop advanced analytics competence and project experience, this challenge can hinder progress for a significant time. ▪ **Regional variation:** Emerging markets, with the exception of China, expressed the most significant concerns. ▪ **Industry variation:** All industries reported they face this challenge, with rough equivalence.	▪ **Impact:** Organizations seeking a broadly applicable AI platform or "an AI" as a solution are stymied by cost and selection paralysis. ▪ **Tactic to overcome:** Begin tactically with AI-powered features on applications and with cloud AI, or with open-source experiments. ▪ **Intensity/relevance/duration:** This challenge will persist for the meaningful duration of AI as a separately considered function or application. ▪ **Regional variation:** Many organizations see more strategic concerns as having priority. ▪ **Industry variation:** All industries reported they face this challenge, with rough equivalence.

ID: 355907

Source: Gartner (July 2018)

- Make the necessity to achieve business value the foundation of any AI project you undertake, and, where possible, favor projects for which AI alone appears to have the potential to deliver that value. Then, because these projects appear suited to AI and nothing else, you can defuse the political question of whether another means would have been preferable. Alternatively, pursue projects that could be solved conventionally, but layer AI on top, or pursue the two approaches simultaneously for comparative purposes.

- Make your CDO or other senior data and analytics leaders pivotal to the formation of any COE or other organizational unit that will pursue AI projects. The CDO should be responsible for ensuring that data is secure and stored in compliance with regulations and policies that are key to the organization's mission, its standing with customers and its ethical positions.

Set Expectations for a Multiyear Arc When Pursuing AI

AI projects employ decision models developed through a combination of evolving tools and technologies. As CIO, you should define a multiyear AI strategy that enables progressive widening of the scope of decisions and use cases, from narrow process applications affecting only a few kinds of customer interaction, for example, to automated self-service interaction routing. Such a strategy should account for unpredictability, due to the complex nature of AI projects. It should include interim milestones to showcase progress. A COE should share demos of outcomes at these milestones with business teams.

AI projects build autonomous decision models based on vast amounts of complex and diverse data. Ensure they choose business use cases that have enough granular data to enable a machine learning tool to operate effectively and improve on analysis conducted via conventional means. Find use cases that can deliver demonstrable business value by solving existing problems iteratively through repeated testing and retraining of decision models.

Select projects that can derive maximum value from AI, such as customer engagement, customer service and business process projects (like supply chain projects). Do so in close collaboration with the business and the CDO to ensure funding doesn't disappear if project timelines drag out. Include the CEO, CDO and line-of-business owners in demonstrations of interim milestone achievements to ensure strong ongoing management support for the project.

Embrace a Complete Strategy, Including a COE, to Mitigate Risk

Returning to the Enterprise Survey, we find that the top challenges to the adoption of AI in the "fear of the unknown" category are:

- Security and privacy concerns
- Potential risks and liabilities
- The question of how to measure the value derived from using AI
- Lack of understanding of what AI is

Mitigate Risk

The more sophisticated the AI application, the greater the expectation for business value — but also the greater the risk of unpredicted failure. CIOs need risk management plans to guide their investments in AI. Weigh up the business value of AI opportunities, your risk tolerance and the cost of keeping AI risk at an acceptable level.

While the potential benefits of AI are great, so are the risks. Inflated expectations and fear of missing out on AI value can lead to decisions that are hasty and poorly thought-out. Examples include opting for a single AI platform, embracing overly ambitious ideas, and allocating AI resources to areas that are starved of the data necessary to render insight. You need patience to help eliminate irrational exuberance, confusion and skepticism.

Your CEO will inevitably ask you about the risks before approving a major investment in AI. To address risk, prioritize AI investments. The temptation may be to use AI to automate tasks and squeeze out every possible cost. But AI applications need human oversight to ensure they deliver value and cause no harm.

Use scenario planning to develop risk management plans. The more complex the AI application, the greater the risk something will go wrong. Develop scenarios of possible failures, how they would affect the business and what should be done to respond in each case.

Address governance impacts by incorporating new regulatory and ethical considerations into your decision making and taking steps to ensure consistency of results from algorithms and models. Foster a data-driven culture and critical data science capabilities to address organizational impacts,

and steer clear of the AI pitfalls associated with selecting technology to measure value.

Identify and Measure the Business Value of Adopting AI

Enterprises are struggling to identify where and how to generate business value with AI. ROI is difficult to measure in this case, for the simple reason that it is too early for many organizations to have achieved it. As is often the case with new technologies, it takes longer to launch AI projects than it does to launch simpler projects with better-known technologies (such as robotic process automation).

Work with your CDO and business executives to develop strategies for data collection, use and impact in order to identify and prioritize the business issues and cases that can best be solved by using AI technologies. AI is particularly vulnerable to misunderstanding, as are other advanced technologies that appear nebulous to business users. Further, its popular profile is so broad and unruly that business executives are inclined to propose AI as a solution for anything, where advanced analysis of conversational techniques is warranted.

Real value, at least initially, is likely to accrue from cost reduction and efficiency more than transformation, if only because these are what organizations are trying to achieve with AI. In the 2018 Gartner Artificial Intelligence Enterprise Perceptions, Plans and Implementation survey, the greatest proportion of respondents — about half — indicated that they will use AI for improved efficiency. This was one of their top three motivations for using AI. The next highest category was cost reduction, which one in three respondents selected among their top three priorities.

Ultimately, we expect to see more transformative and strategic impacts from AI. In the same survey, revenue increase, customer experience improvement and decision making were the next three priorities. Developing a strategic approach to AI will enable organizations initially to address the business value they need, but also to look ahead to use cases that will have greater impact.

Set up an AI COE with a broad view and charter to use AI within business and IT processes. This COE must have clearly defined goals and objectives that address broader business and operational goals.

The AI COE must coexist with and draw on fellow COEs. Any COE devoted to automation or advanced analytics can serve as a foundational contributor to an AI COE — or even as a home for one.

View AI as a strategic lever or an enabling technology to address broader automation, analysis or product development goals. Other potential enabling technologies include conversational agents (such as chatbots), robotic process automation and the Internet of Things.

AI is an enabling technology. A vision for this technology's application could include exploring the use and extent of AI in the organization's operating model or business model in pursuit of a set of measurable business and operational goals. As the organization's familiarity with AI evolves and matures, specific executives should be appointed as sponsors to help the COE fulfill its AI charter, as should leaders for specific programs within it.

The CDO, customer experience officers, security managers and executive leadership must be included in the AI COE. With maturity, it is also expected that organizations will have a strategy that lays out investment plans and that they will commit the funds and resources that are required. Once it reaches a reasonably high level of maturity, the AI practice within an automation COE must demonstrate intent to reduce routine human work throughout the enterprise and foster AI-enabled value-added activities through the actions of executives and AI leaders. These actions should include:

- Making AI part of the enterprise strategy.
- Clarifying the high-level business focus and objectives for AI.
- Promoting augmentation.
- Naming executive sponsors and enterprise leaders for AI programs.
- Designing incentives for AI initiatives that deliver results.
- Committing to fund other fundamental resources required for strategic AI initiatives.

You should also ensure that AI COE leaders formulate a strong governance framework to ensure that decision rights and accountability for acquiring, valuing, creating and using AI intellectual property and artefacts are specified adequately. A governance framework must include principles, guidelines, policies, processes, standards, roles and metrics that ensure AI initiatives will help the enterprise achieve its goals. The governance framework must also align with established enterprise governance components.

An AI COE may nestle within an advanced analytics COE, or exist separately (especially temporarily) in order to attract and incorporate executives and workers from all aspects of an organization. We recommend that you establish goals for COEs that encourage them to inform, persuade, enforce or take responsibility for developing new internal or external products and services. A COE's charter should institute enough flexibility to enable it to change goals and directions over time. The COE should have measurable goals, such as the creation of a system of outreach, registration and recommendation of best practices, or a mandate to infuse AI into internal systems or external products and services. Its purpose should make a clear contribution to the organization's core mission.

An AI COE is important, in part, because of the value of overseeing activities relating to managing, enriching and classifying training data. This must be a role for the COE as it promotes reuse and ensures consistency. All kinds of analysis and incorporation of data into AI projects, from unstructured data mining and enriching to managing ontologies and knowledge graphs, should take place within sight of the COE. It can foster the practice of exploring different learning algorithms which may perform better in particular regional, cultural or business contexts. To gain the skills and knowledge required to avoid obstacles, the COE should collaborate broadly with holders of roles such as security manager, CDO, customer experience specialist, process manager and human resources executive.

Evidence

Enterprise Survey

Results presented are based on a Gartner study conducted to understand the current enterprise technology landscape, with one section focusing specifically on the perspective of senior-level employees. The research was conducted online from November 2017 through December 2017. There were 1,990 respondents from organizations with more than 20 employees located in the U.S., U.K., France, Brazil, China and India. Of that total, 890 respondents qualified to answer the AI section.

To qualify for the AI section, respondents were required to be from an organization that is currently investing in AI (we placed no limitations on the investment). They also had to report that they were at least in the planning stage, and be personally involved any of the following decisions: implementation, planning/budgeting, evaluating vendors, setting strategy.

The results of this survey are representative of the respondent base, but not necessarily of the market as a whole.

The survey was developed collaboratively by a team of Gartner analysts and was reviewed, tested and administered by Gartner's Research Data and Analytics team.

Artificial Intelligence Enterprise Perceptions, Plans and Implementation Survey

Results presented are based on the 2018 Gartner Artificial Intelligence Enterprise Perceptions, Plans and Implementation survey, conducted online in January and February 2018. There were 848 respondents from the U.S. and Canada (208), U.K. (217), China (213) and India (210).

All respondents were screened for active employment in organizations that are piloting or have deployed at least one of the following AI technologies:

- Natural-language processing
- Computer vision
- Artificially intelligent physical robots
- Process augmentation
- Decision augmentation

Respondents were also required to be at least at a managerial level and to have knowledge of the AI budget for 2018. They also had to know about adoption plans for AI solutions and (depending on the AI technology) about strategy, business objectives, business requirements, technology requirements, selection and/or use of providers, effectiveness/ROI measurement, operations management, and/or solution design and implementation.

At the country level, "soft" quotas were established to guarantee a good distribution in terms of AI technology adoption, company size and industry.

The results of this survey are representative of the respondent base, but not necessarily of the market as a whole.

The survey was developed collaboratively by a team of Gartner analysts and was reviewed, tested and administered by Gartner's Research Data and Analytics team.

About Tech Mahindra

Having initiated its practice around the start of the millennium, Tech Mahindra has stepped successfully through the different phases of evolution of BI and Analytics. A global spread spanning various verticals, its core offerings include end-to-end services in Data & Analytics across **13 Verticals, Telecom, BFSI** and **Healthcare & Life Sciences** being amongst the TOP 3. Equipped with **7500+** Data & Analytics consultants with an even mix of Technical, Process and thought-leadership it has several BI Strategy and consulting engagements to its credit across multiple global customers.

The practice includes footprints on **60+ Technologies** with dedicated Data & Analytics Centers of Excellences for end-to-end BI landscape and competencies across all the key technologies. Credited with its impressive credentials of **circa 160 customers**, Tech Mahindra currently serves an active customer base of **135 customers**. With **35+ successful, global BI & Analytics engagements**, Tech Mahindra has consistently been enabling its clients leverage "**enterprise data**" as an asset for growth.

Adopting a differentiated approach, Tech Mahindra has a host of **proven, in-house and certified** BI Solutions that have stood the test of time on addressing the evolving demands of customers in this space. These include offerings across the **core BI & DW, Data Management, Data Science Workbench and AI-Augmented BI Run-Ops**.

Tech Mahindra is positioned as a Niche Player in the 2019 Gartner Magic Quadrant for Data and Analytics Service Providers, Worldwide.

Tech Mahindra represents the connected world, offering innovative and customer-centric information technology experiences, enabling Enterprises, Associates and the Society to Rise™. We are a USD 4.9 billion company with 121,840+ professionals across 90 countries, helping over 935 global customers including Fortune 500 companies. Our convergent, digital, design experiences, innovation platforms and reusable assets connect across a number of technologies to deliver tangible business value and experiences to our stakeholders. Tech Mahindra is the highest ranked Non-U.S. company in the Forbes Global Digital 100 list (2018) and in the Forbes Fab 50 companies in Asia (2018). We are part of the USD 21 billion Mahindra Group that employs more than 200,000 people in over 100 countries. The Group operates in the key industries that drive economic growth, enjoying a leadership position in tractors, utility vehicles, after-market, information technology and vacation ownership.

Contact Us:

For more information contact us at: www.techmahindra.com