Infosys®
Navigate your next

VIEWPOINT

IMPLEMENTING AI IN BUSINESS — CHALLENGES AND RESOLUTIONS

Large enterprises need to implement AI solutions quickly and efficiently, but many don't have the right people, process and technology. Firms can accelerate progress to AI maturity through use of a robust workbench, an operating model built on a data science-led pyramid of skills, and pre-built solutions when data complexity and talent scarcity are too big a challenge.

Infosys® | Knowledge Institute

Artificial intelligence has the potential to cause a revolution in business. The companies that use it best can re-invent themselves and their services, disrupting competition and grabbing market share.

But it's hard to do.

There's a lot of noise and hype around the technology, a cloak that veils the challenge of actually implementing enterprise AI. Only 8% of firms engage in core AI practices that support widespread adoption.[1]

A good business case is often lacking, processing power for running complex algorithms is insufficient, AI terminology keeps executives in the dark and data scientists are in short supply. Many laggard firms don't have the working culture to compete against titans like Google and Facebook, nor do they have the data to create sophisticated AI models. Many turn to providers and platforms to get a leg up on their journey, but this too requires cognizance of the right use case application, and often results simply in "AI for AI's sake."

With the challenges, however, are resolutions that, if carried out effectively, improve efficiencies, create innovative operating models, and — for those brave enough — discover new revenue streams.

What is enterprise AI today?

AI is running Amazon data centers and reading thousands of legal contracts in minutes. It is finding cancerous tumors and predicting when customers will get bored and leave a website. It can identify fraudulent credit card behavior much better than other methods. It can bestow extreme competitive advantage on those that use it properly. Research shows that AI can increase enterprise profit by 38 percentage points, while delivering $14 trillion of additional gross added value to corporations by 2035.[2]

> AI increases profit by 38 percentage points

There are many forms of AI, including image recognition and capabilities that understand natural language. These algorithms process data, understand its context and then take actions based on intelligence.

AI can create meaningful solutions to business problems. Live Enterprise, an Infosys initiative, uses its 'Nia' AI capability to optimize employee and partner requests and resolutions. The platform guides users to take actions that save time, freeing them to do more challenging work.

To get the flywheel spinning, executives must put a strong business case in place and decide which use cases to pursue. They must then understand the implications of this road map on people, process and technology — the challenges they will face, and possible resolutions they can turn to.

AI challenges faced by enterprises

Challenge No. 1: The need for lots of clean data

Google, Amazon, Facebook and Apple (collectively known as GAFA) are good at AI. They have data — lots of data to cleanse for processing. Google can train AI algorithms on the 3.5 billion searches that are made every 24 hours, while the 6 trillion likes posted to Facebook in the same period is a potential gold mine.[3] GAFA didn't set out to be AI companies at their genesis, but the tree of knowledge offers very lucrative fruit that is too tempting not to bite.

Arvind Krishna, IBM's senior vice president of cloud and cognitive software, points out that collecting and cleansing this data is 80% of the work on an AI project, and smaller firms can take over a year to complete this process.[4] Many lose faith and run out of patience.

Challenge No. 2: Complex algorithms

Algorithms exist for processing data, deriving insights and finally taking action on that intelligence.

A retail use case offers food for thought. A given purchase order is fed into the system. The AI needs to uncover whether it is in fact a purchase order as well as its quantity, price, data and other variables. The order number can be triangulated with the vendor supplying the product, providing insights into whether any discrepancies between ordering and delivery were made. This insight might then result in vendor blacklisting if too many discrepancies were found in a six-month period.

This bouquet of algorithms needs to function as a unit and scale up easily to work on enterprise systems. They will vary among industries and will be dependent on data type (a complex document of text and images is very different from a text document of five lines).

Challenge No. 3: Complex data types

Many algorithms exist for processing complex data types such as video and audio. Their accuracy is often questionable, however, and there is the further challenge of scaling the solutions.

Challenge No. 4: The need for faster processing

Complex algorithms that work on huge data sets, such as those found in Amazon Alexa, can take days to train. According to AI consultant Andrew Burgess, any improvement in processor speed will ensure that AI systems remain useful while using up-to-date models.[5] Cloud computing

is a boon here; AI-as-a-service uses specialized hardware that carries out processing on demand — a sure-fire catalyst for the democratization of AI.

Challenge No. 5: Real-time processing

Real-time — or "stream" — processing and generation of insights is challenging without disrupting transaction systems.

"Most of the algorithms today are built on offline insights," says Dr. Raghavan, chief data scientist at Infosys. "The algorithms are scored against a model built on historical data. A new area of research that many experts are looking at is 'How do we build models that are valid for data that was generated in the last half hour?'"

Challenge No. 6: Multiple providers with varying strengths

The sheer variety of AI platforms and vendors can be daunting for any business executive to contend with. There are tech giants such as GAFA that have cloud-based platforms to build models, understand granular data and deploy models into production. Then there are companies like RapidMiner, DataRobot and Dataiku that have built their own frameworks, with features evolving every six months or so. Finally, there are startups exploring just one aspect of the entire ecosystem.

A large enterprise such as a big bank or consumer goods manufacturer, will run proofs of concept to decide the functional areas in which to run these solutions, or whether to use an enterprise wide solution at all. For this, firms must employ in-house talent to understand vendor nuances, the algorithms in use, ease of implementation, data sources, ease of training and so on.

"Evaluating platforms is not easy work. You must test them out, understand their core production and capability, and determine whether they can be used in day-to-day operations," says Dr. Raghavan.

Challenge No. 7: Difficulty of knowledge sharing and code reuse

In larger firms, different divisions may develop their own algorithms, leading to duplication if code is not shared between teams. These inefficiencies increase costs and ultimately lead to inconsistency across the enterprise. This can even affect how regulators view a firm's use of AI, particularly if it seems that there is no single methodology or framework in place.

Challenge No. 8: Dependence on niche, expert talent

There is often a big talent gap between firms already moving on their AI path and those still evaluating the technology. Those without momentum often give up due to lack of dedicated resources.[6] One reason is the mammoth gulf between the number of AI practitioners (300,000 worldwide) and the number of job postings (in the millions). The best candidates go to companies like DeepMind, Google, Airbnb and Intel — and others that tout exciting use cases and significant monetary rewards.[7]

> The 300,000 AI practitioners worldwide can choose from millions of AI job postings

Even more crucial is finding talent with industry-specific knowledge. If found, these workers become indispensable to a project until its completion.

Challenge No. 9: Developing the products

Once AI algorithms are in place, the business needs to commercialize the opportunity through the creation of products and services. This requires significant input from product designers, business analysts and business development teams. Those involved in the service design need to pay particular attention to the interfaces and user experience, and for many designers, AI is an area where they will not have expertise.

Embedding complex algorithms into existing operations and automation is also important. This calls for creating production grade AI applications that will work in legacy systems. A large telco such as AT&T might be dealing with 130 off-the-shelf software products, with data coming from 60% of those systems, which in turn may come from 10,000 databases — not an easy task for any production shop.

Challenge No. 10: Stakeholder involvement

"Getting business and IT stakeholders around a table is a significant challenge at most firms," says consultant and author Andrew Burgess.[8]

So too is managing expectations of business executives once they do collaborate with IT. For many, AI is seen as a magic pill that will reduce the costs of operations while improving production quality.

"Many executives don't realize that it's complex to create AI applications," says Dr. Raghavan. "Often, executives have unrealistic estimates of production, hoping that a new system will be ready in the next two to three weeks."

Resolutions

A robust 'workbench' approach can be used to circumvent many of these challenges. Firms should also build up capabilities across people, process and technologies. Pre-built solutions are a way forward for those challenged to secure top talent or that have an AI-averse culture.

The robust workbench

A workbench brings together development, deployment and ongoing maintenance under one roof. It offers users a factory-like

process, with automation and DevOps embedded in the process.

"The AI workbench encapsulates the data layer for a data scientist," says Dr. Raghavan. "All they need to worry about is where to pull the data from, giving them freedom to build and validate models across geographies where hundreds of scientists will be working on the same problem."

This means that models can be standardized, and algorithms can be reused on a global scale. It also provides a single access point for validated models to be deployed into production. The workbench should consist of the following elements:

- *API-driven development*

A good workbench offers Application Programming Interface-driven development, which forgoes dependence on ultra-niche talent. It also enables code reuse and knowledge sharing. Here, interactions with a program are standardized and essentially static. For example, IBM's Watson (the heavily lauded AI capability) is just a series of APIs that each carry out a specific function, such as speech recognition or Q&A, which can be called by another program with the right access. APIs also democratize AI, since the value is no longer in the algorithm but in the amount of training that is carried out on a given dataset.

- *Containers*

Containers — encapsulated applications with their own operating system and memory — mean that an AI application can be published as a service and can be maintained and monitored very easily. From a deployment perspective, containers ensure that downstream applications don't worry about specific aspects of modules, libraries and functions or the AI application itself. Containers help with code reuse and sharing (like API calls) and reduce the number of complex algorithms in use.

- *Multicloud API enabled*

Even within one enterprise, a good workbench must have the flexibility to support multiple cloud environments, in order to foster innovation and preserve business unit autonomy (the finance division in the U.K. might use Microsoft Azure, for example, while procurement in Germany uses Amazon Web Services). This in turn solves the challenge of getting business and IT stakeholders around the table to push the AI strategy through quickly.

> A good AI workbench must have multi-cloud flexibility to foster innovation and preserve autonomy

- *Guided experience*

The guided experience is a "playbook" that increases momentum in any AI project during the design stage. The point is to showcase how the AI system will work once deployed, which can be used to attract data scientists to the project and ensure the right partner comes along for the journey. A retail banking chatbot can be simulated, for instance, offering guidance to key stakeholders on its look and feel, implications for back-end data sources and ultimately the user experience.

The guided experience is crucial for AI projects with complex algorithms and architectures. If solution architects overlook this part of the workbench, project development degenerates into "spaghetti," with too many what-ifs occurring in design.

Capability buildup

- *Data scientist-led talent*

Very few firms make the most of AI capabilities, which, with some thought, can be integrated through effective partnering and ready-made solutions. Most often this is due to a lack of in-house data science skills and AI expertise that prevents seeing the forest for the trees. Building up a pyramid of skills dedicated to AI-specific systems is crucial.

For instance, data scientists classically use structured data. However, with the proliferation of text, speech, image and video, an organization will need to build up capabilities in those specific areas too. Once the skills pyramid is created (a bank might need connoisseurs of text data while a media firm will need video and image analytics expertise), firms will be more confident that skilled workers won't leave them partway through a project. They will also find that their reputation as an AI-driven organization will be enhanced, giving them the motivation to meet any one of the other 10 challenges we listed.

For firms that are not highly technical or those that are less mature, we advocate hiring experienced professionals who can evaluate external talent that will ultimately build up the pyramid.

- *Playbooks for specific knowledge*

Here, documentation and "wiki pages" for a particular project are developed to detail certain aspects of use cases that an enterprise wants to build out. For instance, a telco or retail company might develop a recommendation engine. The "next best action" of the engine will be built on playbooks that detail what specific components need to be created and what data sources need to be included. This reduces the complexity of building the right code and simplifies dialogue between business and IT CXOs.

- *Provider-specific training and hands-on experience*

If an enterprise wants to build an Azure stack, an Azure developer is needed who knows the Azure data lake and is familiar with Azure machine learning, Azure Cortana and so on. This might seem obvious, but is often overlooked. So too is ensuring that workers have used a particular platform, in

a particular style and in a particular industry. If predictive AI for a mining company is the use case, having someone familiar with the mining industry will ensure that they solve the right problems, since they talk the same language and understand the domain. Industry-specific experience also helps ease adoption. The relevant person will know exactly where to procure a platform or pre-built solution, and will know how to influence key stakeholders, while enabling easier adoption of AI into the overall system.

Pre-built solutions

Access to AI software libraries speeds up development and deployment, while proof-of-concept (POC) solutions ensure that the product is built cheaply and quickly. Pre-trained models working on huge data sets help democratize AI and counter the talent challenge.

- *A pre-built box of algorithms*

In-house development involves a lot of cost; identifying the right resources to create AI algorithms is expensive both in planning and development. Having a pre-built library is a much more convenient way of developing an algorithm that works for any given use case.

"If, for example, I'm doing clustering or time-series regression, having a library is always a blessing," says Dr. Raghavan.

If artifacts are available at the library level, then a developer can custom build any application very quickly.

Sentiment analysis is a good example of a pre-built box of algorithms to solve problems. Sentiment analysis can be embedded as part of a corporate communications application, or it can be part of a product development application. Here, pre-built components parse text. There are at least 10 different steps in the analysis, starting with speech detection and moving on to sentence detection, sentence bonding and looking at correlations across sentences. Each of these steps use standard elements available in open source.

"You can then wire step one to 10 as part of the sentiment analysis and create a pre-built API," says Dr. Raghavan. "Anyone who wants to use that API might have to do some minor customization, but they won't build anything from scratch."

So when it comes to sentiment analysis, building and deploying solutions becomes much easier if that pre-built layer already exists.

- *POC solutions*

POC solutions are part of a natural agile cycle for building and deploying AI solutions, allowing firms to get critical insight into AI platforms, potential providers and ensuring that the product is built cheaply and quickly.

To build an AI solution, underlying challenges must first be understood. Very rarely is the first solution similar to the final solution. Different algorithms must be tested, along with different software vendors and software frameworks, and the solution must be tested in the user environment. The POC solution created builds confidence and, more importantly, ensures that the project runs on time and isn't a failure. Once complete, the solution is scaled up.

- *AI-off-the-shelf*

Infosys has built a pre-trained functionality known as Document Comprehension.

"When you upload a document, however complex, the pre-trained model is so sophisticated that it can derive questions and answers from the image and document by itself," says Harinder Cour, a consultant at the Infosys Center for Emerging Technologies. "If you don't like those questions or answers, you can ask your own questions and it will scan the entire document in real-time and give you a good answer nine times out of 10. In the future, insights will be derived from even more complex data types such as images and video."

This pre-built solution effectively reduces the dependence on expert coders and solves the problem of processing complex data types in real-time — a truly disruptive technology if ever there was one.

The future enterprise

AI is guiding decisions on everything from bank loans to crop harvests. And it's big business. Harvard Business Review estimates that AI will add $13 trillion to the global economy over the next decade.[9] However, as of this writing, most firms have only applied AI in a single business process. To improve efficiencies and identify new revenue opportunities, many challenges must be resolved, including scarcity of talent, risk-averse cultures, and an inability to imagine what the solution will look like on completion.

We advocate starting with a workbench that bridges the gulf between development and production. Building capabilities in-house must also be a priority, accelerating changes across people, process and technology. Finally, if these skills are hard to come by, pre-built solutions trained on big data can be employed to generate momentum toward AI maturity.

Before thinking about implementation, executives must put a strong business case in place, and decide which use cases to pursue. They must then understand the implications of this road map on their particular industry. Only then will they have clarity to make it back to shore and turn data into cash.

References

1. Building the AI-Powered Organization, HBR, 2019
2. How AI Boosts Industry Profits and Innovation, Accenture
3. Andrew Burgess, The Executive Guide to Artificial Intelligence, Palgrave Macmillan
4. Data challenges are Halting AI Projects, IBM Executive Says
5. Andrew Burgess, The Executive Guide to Artificial Intelligence, Palgrave Macmillan
6. AI Adoption Held Back by Company Culture, Talent Shortage, Data Issues", WSJ, 2019
7. It's Recruiting Season for AI's Top Talent, and Things Are Getting a Little Zany", MIT Technology Review
8. Andrew Burgess, The Executive Guide to Artificial Intelligence, Palgrave Macmillan
9. Building the AI-Powered Organization, HBR, 2019

SME

Dr N.R. Srinivasa Raghavan

Global Chief Data Scientist
Infosys

Authors

Harry Keir Hughes

Senior Consultant
Infosys Knowledge Institute

Isaac LaBauve

Senior Consultant
Infosys Knowledge Institute

About Infosys Knowledge Institute

The Infosys Knowledge Institute helps industry leaders develop a deeper understanding of business and technology trends through compelling thought leadership. Our researchers and subject matter experts provide a fact base that aids decision making on critical business and technology issues.
To view our research, visit Infosys Knowledge Institute at infosys.com/IKI

For more information, contact askus@infosys.com

Infosys®
Navigate your next