IBM Services
White Paper

Beyond the hype: A guide to understanding and successfully implementing artificial intelligence within your business

Contents

Introduction

To implement AI within your organization successfully you need to understand what AI is, where it currently stands, what value can provide to businesses and how it can be successfully adopted. This white paper is written for business leaders looking for practical advice on how to leverage artificial intelligence (AI) for their organizations.

The fact that AI has been hyped doesn't take away from its capabilities as a real value driver. Heavy investments have been made in AI across multiple industries; the Chinese government even made it one of their core sectors. And companies like IBM, Microsoft, Google and Amazon are leading the pack when it comes to utilizing data and AI.

AI provides enormous amounts of value in multiple industries. Because of its high value potential, many companies have been scrambling to implement AI within their organizations. And the projects, when implemented properly, have shown significant returns and improved competitive edge. If your company hasn't started implementing AI, it may lag behind its competitors, so it's critical to evaluate what AI can do for your organization. But this doesn't mean you should hire data scientists or acquiring data science solutions without a clear strategy. Implementing AI should be a carefully thought out process. Otherwise it may turn out to be a costly failure.

The authors of this paper want to provide you with the knowledge you need to evaluate what type of AI solutions you can implement to give your company a competitive advantage.

What is AI?

When people hear AI they often think about sentient robots and magic boxes. AI today is much more mundane and simple—but that doesn't mean it's not powerful. Another misconception is that high-profile research projects can be applied directly to any business situation. AI done right can create extreme return on investments (ROIs)—for instance through automation or precise prediction. But it does take thought, time and proper implementation. We have seen that success and value generated by AI projects is increased when there is a grounded understanding and expectation of what the technology can deliver from the C-suite down.

We are now at the brink of the fourth Industrial Revolution. AI is one of the biggest facets of this revolution, and it will affect almost all sectors, as did previous Industrial Revolutions.[1] AI's abilities have increased significantly since its inception in 1955; it can now detect patterns more accurately, continuously and based on more data. Currently, AI has surpassed human intelligence in some specific domains. These domains can be split in three categories: general tasks, formal tasks and expert tasks. General tasks could include visual recognition, speech recognition, natural language processing and translation. Formal tasks are related to games where some theorem and learning is involved. And expert tasks are those that would otherwise be executed by a domain expert. (Think of tasks such as diagnosing disease and engineering.[2])

As already stated, AI was first named in 1955 and was defined as the ability of machines to perform human-like tasks. The term has gained popularity ever since its first mention. However, there is still quite a bit of confusion about the difference between AI, machine learning and deep learning—but simply stated, AI encompasses the latter two.

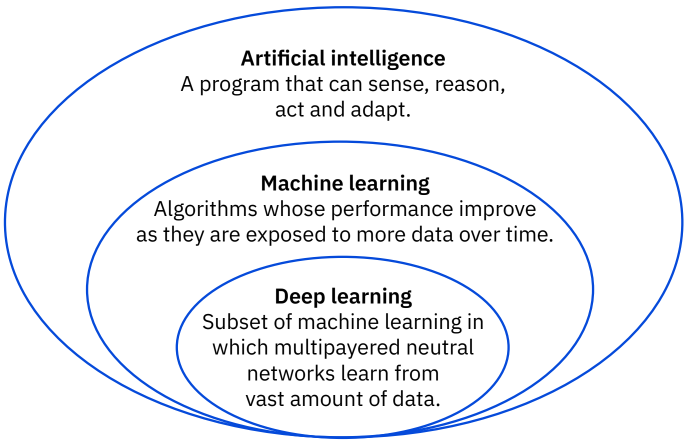

Figure 1: The distinction between AI, Machine Learning and Deep Learning.

AI

"Artificial Intelligence (AI) is a science and a set of computational technologies that are inspired by—but typically operate quite differently from—the ways people use their nervous systems and bodies to sense, learn, reason and take action."[3] Lately there has been a big rise in the day-to-day use of machines powered by AI. These machines are wired using cross-disciplinary approaches based on mathematics, computer science, statistics, psychology and more.[4] Virtual assistants are becoming more common, most of the web shops predict your purchases, many companies make use of chatbots in their customer service and many companies use algorithms to detect fraud. These are just a few of the examples of how AI is used every day.

Machine learning

Machine learning is enabling a machine to learn from data without explicitly programming it with rules, because it can learn from the data it's given. In essence, you could build an AI consisting of many different rules and it would also be able to be AI. But instead of programming all the rules, you feed the algorithm data and let the algorithm adjust itself to improve the accuracy of the algorithm. Traditional science algorithms mainly process, whereas machine learning is about applying an algorithm to fit a model to the data. Examples of machine-learning algorithms that are used a lot and that you might be familiar with are decision trees, random forest, Bayesian networks, K-mean clustering, neural networks, regression, artificial neural networks, deep learning and reinforcement learning. Artificial neural networks and deep learning have recently become more common machine learning algorithms.

Implementation examples would be predicting stock market prices or predicting whether a customer will churn from your company.

Deep learning

Deep learning (DL) is a relatively new set of methods that is changing machine learning in fundamental ways. DL isn't an algorithm per se, but rather a family of algorithms that implements deep networks (many layers). These networks are so deep that new methods of computation, such as graphics processing units (GPUs), are required to train them, in addition to clusters of compute nodes.

DL works very well with large amounts data, and whenever a problem is too complex to understand and engineer features (due to unstructured data, for instance). DL almost always outperforms the other types of algorithms when it comes to image classification, natural language processing and speech recognition. An example would be recognizing melanoma or conducting machine translation, which was not possible using previous techniques.

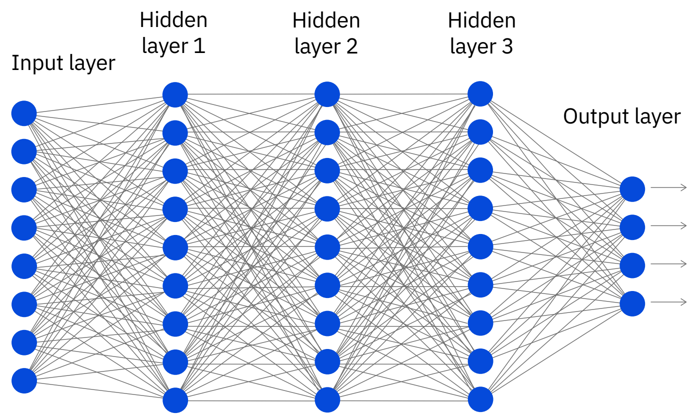

Figure 2: Deep neural network with five layers.[5]

Currently, the larger the neural network and the more data that can be added to it, the better the performance a neural network can provide. DL is very powerful, but it has a couple of drawbacks. It's almost impossible to determine why the system came to a certain conclusion. This is called the "black box" problem, though there are now many available techniques that can increase insights in the inner workings of the DL model. Also, deep learning often requires extensive training times, a lot of data and specific hardware requirements, and it's not easy to acquire the specific skills needed to develop a new DL solution to a problem.

In conclusion, there is no one algorithm that can fit or solve all problems. Success really depends on the problem you need to solve and the data you have available. Sometimes a problem will need a hybrid approach, where you use multiple algorithms to solve the problem. Each problem requires extensive investigation of what constitutes a best-fit type of algorithm. You should take into account transparency and how much data, capabilities and time you have, because some algorithms take a long time to run.

How does an AI system learn?

To illustrate how AI system learning works, we'll next describe what a data scientist does, and what a machine does in the process of developing AI solutions. Later, when we discuss the pitfalls of implementing AI, we'll explain what types of skills you'll need to successfully build a data science team.

Data scientist

The data scientist extracts knowledge and interprets data by using the right tools and statistical methods[6]. The data scientist first helps identify data-analytics problems. Next, he or she defines defines the right algorithm and tools. He or she cleans the data and collects the correct data for the problem. The next step is to define hyperparameters and engineer the features in a way that it fits the model. After the model provides output the data scientist analyzes and identifies patterns and trends. Then the data scientist communicates the results to the stakeholders.

Machine

The machine learns to recognize patterns in the data that it is fed to it, and then maps these patterns to future outcomes. The machine learns through adjusting the weights and biases in the network from feedback to get to the correct outcome. This feedback must come from a trainer—the data scientist. The data scientist tells the model what should happen and what shouldn't happen. This correction is then sent back through the network and an error rate is computed. With each iteration, the model works to decrease the error rate.

In the appendix section of this paper, we have a section that addresses the different types of algorithms' learning. There are four main types of learning; supervised, unsupervised, transferred and reinforcement learning. While unsupervised, reinforcement and transfer learning show great potential, supervised learning is currently the type that can provide the highest economic value.

What has driven the development of AI?

With increasing computing power and more data, the potential value of algorithms became higher. People and companies saw possibilities to embed these smart systems into their companies to drive strategy and innovation. As the power of algorithms, computing and amounts of data increased companies started to see an increasing amount of use cases. AI started to become an essential value. Companies saw that these systems could move them closer to their customers, drive efficiency, enhance employee experience and capability, automatize tasks, decrease costs and improve revenue.

But AI went through some lulls and spikes before it became capable of enabling so many benefits. At a certain point the general public gave up on AI and machine learning as a whole, which stalled developments and investments. For instance, when Defense Advanced Research Projects Agency (DARPA) worked together with Carnegie Mellon to implement speech recognition for their pilots, they cut the project in 1974 after having spent millions of dollars on it. From the 1960s to 1974, the main funding for AI came from governments. But after 1974 there was barely any government funding because of multiple failed AI projects, so the belief in the feasibility of AI declined. But after this period, improvements in efficiency lead to successful business cases, once again proving value. Today we now see AI as one of the big value drivers for companies, and to compete most companies must adopt AI strategies. Three topics have made AI available for many companies right now:

– **The evolution of data:** A factor contributing to the massive adoption of AI is the exponential growth of available data. With the introduction of the Internet, social media, proliferation of sensors and smart devices, and the fact that data storage became cheaper, it became more accessible than ever before.
– **The evolution of algorithms:** Algorithms have been around since we could write. Recently, the development of more advanced algorithms has helped AI become more powerful and efficient.
– **The evolution of computing:** Another major factor in AI's current success is its computing power. Back when AI was just beginning to be developed, the computing power was minimal. Computers nowadays can take much more data and heavier algorithms than in the 1950s.

These developments would not have taken place without significant investments and proven business value.

1950	Turing publishes the Turing-test. "The point at which a machine has answers like a human"
1955	AI first named by John McCarthy
1956	"First" AI algorithm Logic Theorist by Simon and Newell
1957	Rosenblatt invents the first self learning algorithm with the perceptron
1958	IBM 305, the first hard drive, 5 MB
1969	Backpropagation, one of the most important areas of a neural network, is proposed
1970	IBM 1330, 100MB
1974	Intel produces second generation general purpose chips
1974-1980	First AI winter, the belief in machine learning and AI had dropped after multiple unsuccessful experiments combined with insufficient computing power, network capabilities and database capacity
1985	IBM 0665 hard drive, 40 MB. But much smaller than the 1330
1989	First convolutional neural network developed (used a lot in image recognition)
1991	The internet is open for the public
1992	First versions of natural language solutions set up.
1997	IBM's deep blue defeats Kasparov in Chess
1998	Google's Page rank is published
2000	The adoption of Internet by the masses
2002	Amazon brings cloud computing to the masses
2004	Google develops an algorithm that can handle large amounts of data faster.
2005	Stanford Robot drives automatically
2006	IBM introduces Watson. A question answering machine that later wins from a Jeopardy champion
2010	Worldwide IP traffic exceeds 20 exabytes (20 billion gigabytes) per month
2012	Facebook gets a billion users
2014	There are more mobile devices than humans in the world
2018	Project debater of IBM shows ability to process very large data sets, including millions of news articles across dozens of subjects, and then turn snippets of arguments into full flowing prose—a challenging task for a computer

Figure 3: AI timeline.[7]

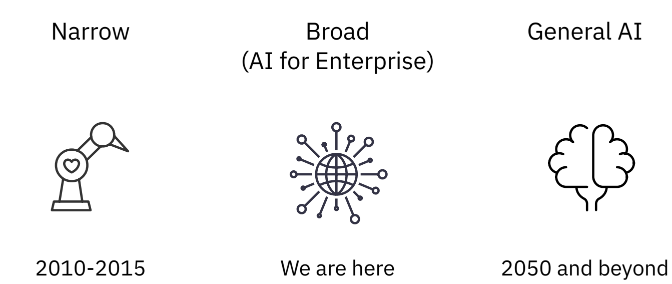

Figure 4: The three categories of AI.

Where are we today with AI?

We can split the term AI into three categories: general, broad and narrow. General AI encompasses all the human-like capabilities, whereas narrow AI can only do a certain task—and it can do it quite well—but narrow AI can't transfer its knowledge to different sorts of problems.

Narrow AI
Narrow AI is focused on addressing very focused tasks (such as buying a book with a voice-based device) based on "common knowledge." That's the reason narrow AI is scaling very quickly in the consumer world where there are a lot of common tasks and data to train these systems. Narrow or weak AI is, contrary to the naming, very powerful at routine jobs.

Broad AI
What we see today in self driving cars is still defined as narrow AI. You can see it as a collection of narrow AI systems that make decisions. This is what we call broad AI. Another example of broad AI includes a system within a bank that analyzes the balance sheet of corporate customers to recommend the best currency hedging strategy. Another example would be a system that supports engineers who work on complex maintenance tasks on a platform in the middle of the Atlantic Ocean. Broad AI is about integrating AI within a specific business process of an enterprise where you need business- and enterprise-specific knowledge and data to train this type of system. These tasks are very different from the narrow AI used in the consumer world because the data and knowledge available in the enterprise are much more limited in terms of volumes, very industry specific and in most of the cases private (for example owned by an enterprise). This is what we believe is currently the most valuable type of AI currently for the enterprise.

General AI
General AI is far from reaching its potential. The expectations are that it will take at least another couple of decades. General AI refers to machines that can perform any intellectual task a human can. Currently AI does not have the ability to think abstractly, strategize and use previous experiences to come up with new creative ideas as humans do.

Some people think we will have general AI in a couple of decades others like IBM's Rob High and Google's Peter Norvig believe we don't need broad AI at all.[8]

What are the areas in which AI provides the most value today?

While relevant AI use cases span various areas across virtually every industry, there are three main macro domains that continue to drive the adoption as well as the most economies across businesses. They are:

- Cognitive engagement: Involves how to deliver new ways for humans to engage with machines, moving from pure digital experiences (such as the ability to run transactions digitally) into human-like natural conversations.
- Cognitive insights and knowledge: Addresses how to augment humans who are overwhelmed with information and knowledge.
- Cognitive automation: Relates to how to move from process automation to mimicking human intelligence to facilitate complex and knowledge-intense business decisions.

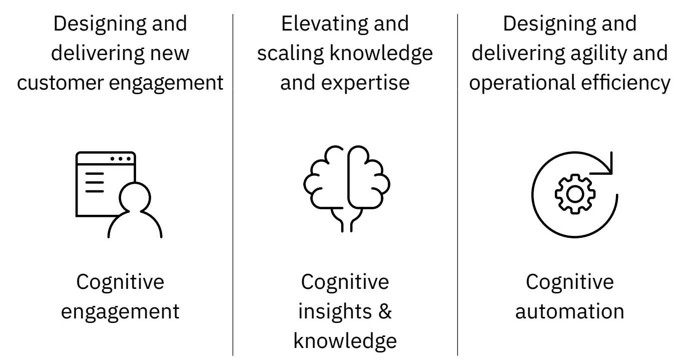

Figure 5: The most valuable AI implementations.

What are some examples of successful implementations?

There's a vast amount of problems AI is already addressing to deliver business value across the three macro domains described in the previous section. We want to explain a couple of use cases our IBM team has successfully completed to demonstrate where AI can bring value.

Manufacturer: engine anomaly detection using a neural network

Using the many different available sensor measurements from large truck engines, a neural network was trained to recognize normal and abnormal engine behavior. In the huge, high dimensional (many variables) dataset the neural network learned the natural correlations and relationships between all different readings. The resulting model was able to predict "normal" values given certain operating conditions and could thus also be used to detect when specific measurements were out of the ordinary. Anomalous sensor readings are highly predictive of pending engine failures.

Car manufacturer: predictive failure detection for welding robots and predictive maintenance assessment

Through supervised learning techniques, predictive models were developed that could provide an early warning of failure based on the different system messages and sensor readings that continuously stream from the production line. This early warning could be used to prioritize maintenance and reduce both downtime as well as false positives and needless efforts. Working through this first proof of value, the data scientists uncovered many data quality challenges that could be worked around to realize more value.

Utility company: micro-grid energy forecasting and production mix optimization
The output of machine learning-based predictive models with prescriptive, mathematical optimization models to prescribe the optimal mix of power production sources to meet predicted demand and to minimize costs. This required both the prediction of demand as well as prediction of available solar and wind energy capacity.

Material producer: insights dashboard
IBM worked with the client's sourcing experts to understand the business dynamics and create inventory of possibly relevant data sources. Several machine learning models were then trained to learn the price behavior and forecast future price development. The models also enabled buyers to evaluate their own "what if" scenarios. This was further supplemented with IBM® Watson Discovery News service, which identified the most relevant news articles related to the material of interest. This all came together for the user in an interactive dashboard to consume the insights and interact with the data and models to make buying decisions.

Best practices to successfully implement AI within an organization

Let's discuss what you need to do before implementing AI. Currently, many companies are scrambling to implement AI within their environments because they believe it will keep them ahead of the game—which, if thought through, is the case. But here are a few steps you need to take.

We've defined three main steps to implement AI in your company. These steps are:

- Develop an AI strategy and roadmap
- Establish AI capabilities and skills
- Start small and scale quickly

Develop an AI strategy and roadmap
First, it's important to understand AI and to research what it can and can't do for your organization. You can get more familiar with AI by collaborating with a data scientist, because it's important that the C-level has a good understanding of AI and its implementation difficulties before they define where and how to implement. What often happens if AI is not holistically understood the overall project won't provide value.

Once AI is understood, the next question you should ask yourself is: "What specific problem do I want to solve, or what opportunity do I want to take?" Is your company looking to drive efficiency in the back office, differentiate its digital proposition, generate new revenue streams by leveraging customers' insights or even reinvent its entire business?

After having thought this through, you'll probably have many different use cases. At this point, it's critical to prioritize these cases into a transformation roadmap that covers both a long-term vision as well as concrete feasible quick wins. Next, you should think about what data you have available. To solve most of the problems with AI you need to have relevant data. Without data, AI will not provide any value. For many companies it's a task in and of itself to keep track of the type of data—as well as where it's stored and in what way. Often the first step will be to understand the data you currently have and the type you need to implement your AI case.

Establish AI capabilities and skills

AI requires a completely new set of capabilities and skills which may be in short supply in your organization. To build the required in-house AI skills, it's important to plan, establish and grow a dedicated Center of Competence or leverage the IBM Garage concept to perform in partnership. Not only this dedicated team is important, but you also need to assure the right mindset and way of working in the rest of the organization. It's critical that these functions occur in conjunction with developing and integrating an AI platform within your current IT architecture to implement and scale AI.

Start small and scale quickly

– Start with minimal valuable products (MVPs)
  In this phase you want to bring in experts to help quickly develop solutions to your business problems. This can only be done once the before-mentioned steps are completed and the business is ready organizationally and technologically. This also means that the experts you bring in should be both business and technologically savvy. A good duration for a MVP is normally between two and three months. Our experience shows that starting with large-scale, complex and very long AI implementation projects normally lead to failure.
– Set understandable key performance indicators (KPIs)
  To make sure that a project will succeed, you need to define KPIs that are understandable for your business—including employees and other stakeholders. These KPIs will help you evaluate whether a project is successful. In general, we suggest taking a second look at these KPIs after an appropriate duration to decide whether the project is successful or if you should discontinue it. If your business can't pinpoint the right KPIs to measure success, the project is too complex.
– Roll-out through company (culture)
  Once agreement is reached about which projects would be worth working on, it's time to implement the MVP within your company. It's important that the way you implement it is looked at from both the business and the technical side.

Pitfalls of AI implementation

With the experience of implementing many of these cases we normally see a couple of problems organizations have when implementing AI. We've listed them because we feel it's important to know what you should consider when implementing AI.

Culture
Pitfalls: Looking at the advances we've made over the last decade, gathering data is easy. But it's what's done with the data that provides the most value. The biggest pitfall we often encounter is a culture that's not committed to making data-driven decisions. Examples are cultures that can't innovate in an agile fashion or can't leave room for trial and error, as are cultures that have traditionally been unwilling to transform a process. This reluctance usually has to do with fear of job loss, skepticism or a knowledge gap. It's a challenge to get your organization ready to embrace a data-driven culture, and for many of your employees it can feel like a counter intuitive process.

Recommendation: Digital change management, training and preparation for the shift in thinking is required. You should start with small wins that are visible for relevant departments. The business users should be your starting point for agile development and in your design thinking process. All the results that come out of the project should be measurable; this way you can easily show your wins. Then, for a data-driven culture to take hold, the whole organization must embrace it. The message must be clear for all employees: "Decision are made based on data."

Building trust within the company
Pitfalls: This step is often overlooked, but when you want to implement AI you need get project stakeholders involved and on board. In Design Thinking, IBM starts with the business and its users. A data scientist will require domain knowledge and access to data, and the stakeholders should accommodate this need to help to speed up the process. Also, consider the need for AI education and devote time to considering the right user interface.

Recommendation: Use a form of change management to establish user adoption. IBM's best practice is to leverage its digital change approach to involve the users in the development of the project through studio's in which we envision and co-create together. AI is not something everyone will be able to grasp immediately. Working with AI should therefore be carefully implemented in the business environment.

Expectation management
Pitfalls: In many cases we work on, we see that employees or other stakeholders don't believe in AI or think it's a magical box we take from our office that will quickly solve all the problems a company. This leads to disappointment when an AI project is not delivered in a short timeframe or if it doesn't deliver the expected results. And this disappointment can lead to a lack of belief that will eventually diminish the will to implement AI and experience the benefits of its long-term possibilities.

Recommendation: Start with developing a solid strategy and roadmap. Define where you need to go and what you need to add to the organization to get there. First steps often include mundane items such as data governance and warehousing, but you need to take these steps to properly implement AI. It's therefore important to implement proper change management to guide the company.

Bad data
Pitfalls: One problem we often come across is bad-quality data. This problem derails, limits and complicates many machine learning and AI projects. Bad or “dirty”data can mean fields are missing, that there are duplicates in the data, or that it’s outdated data, contains spelling or punctuation errors and is generally incorrect. As we are moving toward more data-driven decision making in enterprises, it’s absolutely essential to have clean data. Outcomes derived from bad data will lead to incorrect decision making.

Recommendation: To overcome incorrect decision making based on bad data we suggest the enterprise incorporate standardization of data, monitoring of data, cleaning of incoming data and a centralized control of data.

Sponsorship
Pitfalls: When implementing an AI strategy, it’s important that the right people, such as department heads, CxOs or managers support the project. Implementing AI through the organization can be a long process, and without support of the right people there’s a higher probability that the project will fail. What you might see when this problem occurs is that your employees might not put enough time into problem definition and subject knowledge sharing.

Recommendation: Align the right people before you start in a “garage” concept. The key stakeholders should be identified and should bring input and willingness to the table. Create buy-in and support with employees, other stakeholders, management and C-suite.

Lack of capabilities
Pitfalls: Many companies want to implement machine learning right off the bat. But to implement AI, you need to look at two things: the first is acquiring or outsourcing your own data science talent, and the second is looking at your current IT infrastructure.

If you choose the in-house route, you’re taking a more challenging route. This option can be very rewarding but it’s important to take in account that it will require more time to set up infrastructures, pipelines and research. If you decide to take this route, you’ll need to acquire the following skills in house:

- Researchers to create new solutions to your products
- Project managers to keep the team on track
- Domain experts who have knowledge about your products, customers and the business environment surrounding the product
- Data engineers and machine learning engineers who can scale the algorithms
- Data analysts who can process the outcome
- Statisticians to help ensure quality results
- Software engineers to turn all you’ve created into something that can be used by the masses—be it your customers or your employees

While these roles don’t all need to be filled by individual employees, it’s essential to have all these skills in house.

The second option is easier and quicker to implement. In this case you’d make use of the capabilities of an external party such as through the IBM Garage offering, which can give you the ability to use these capabilities without having to set up a complete internal department.

Recommendation: Before starting an AI project, you should evaluate the capabilities you currently have inhouse, and then define the cost you'll occur acquiring the capabilities needed to close any gaps. Then you can define whether you want to hire these capabilities internally or use an external experienced resource to access the required capabilities.

Scalability
Pitfalls: To properly scale the correct architecture, integration and employees that know how to use them need be in place. Many data scientists believe that the research and development (R&D) of a data science project is similar to a scaled IT implementation, yet the two are very different. Also, because there is increasing enterprise demand for AI, organizations want to analyze large amounts of data. This can require quite a bit of time to train an algorithm—days or even weeks depending on the amount of data involved. We still encounter data scientists trying to perform this function on their laptops.[9]

Recommendation: Develop a plan to set up the correct AI architecture, a platform to deploy to, a data integration strategy and properly trained data scientists.

Not enough (available) data
Pitfalls: Data availability depends on the company and on how it stores data. Some organizations have the data but don't have it readily available. Large corporations are challenged to locate and keep track of the right data. Smaller companies may be challenged by the amount of data they produce.

Recommendation: Not all problems need machine learning and AI. If you find yourself in a situation where you don't have enough data, you should carefully consider if you should launch your product using machine learning. If you choose that option but don't have enough data, you can overcome the problem by acquiring external data or by using simpler models.

Unlabeled data
Pitfalls: A common problem when implementing AI is data that's not classified by humans or machines, which leaves you unable to train the system (in the case of a supervised algorithm). For instance, if you wanted to predict fraud, but the historical fraudulent cases weren't labelled, it would be impossible for the algorithm to map input to outputs.

Recommendation: In this case, define what data you'll need to run a sound model. This data will then need to be labeled—which can be expensive—or you can use algorithms to accomplish this task. It's a more advanced technique, but you can use reinforcement learning or semi-supervised models. Your data scientist might know which approach is best.

Explainable results
Pitfalls: In many cases, the business wants and needs to know why specific outcomes occur. This is in line with one of IBMs principles for ethical use of AI is: AI systems must be transparent and explainable. Depending on what kind of algorithm is used, an AI can't show exactly which variables brought it to a given conclusion. We call these black-box algorithms. While they sometimes may perform better, they don't give a lot of explanation. For example, you might have an algorithm that defines whether an applicant gets a loan. According to European law in some countries, if you deny the applicant, you need to explain why you did so.

Recommendation: There are several new algorithms that help explain what happens in the black-box models such as Local Interpretable Model-Agnostic (LIME).[10] But this doesn't solve the problem in all cases. You should therefore define what you're going to use the model for. If the model needs to explain precisely why it came to a certain conclusion, your data scientists should consider other models that provide the specific information required. Ethics and AI should go hand-in-hand.

Summary

AI has the potential to bring a lot of value to your company if thought through and implemented properly. The authors of this paper hope we've made it clear how you can achieve this goal. We've discussed what AI is what it can do for your organization, how it should be implemented and what pitfalls you should avoid once you've decided to implement AI.

To sum up this paper's key points:

AI not something of the future, it is real today, and it fuels the fourth Industrial Revolution: As you will see later in this paper, there are many cases where AI is being successfully implemented and driving competitive advantage. Companies such as IBM, Nvidia, Twitter, Delta Airlines, Walmart, Netflix, Spotify and Kreditech show that their data-driven approach produces extremely valuable business models. But for many companies, failing to implement a data-driven strategy can lead to lost market share.

Properly implementing AI requires careful evaluation and planning: You need to evaluate how AI can help solve your problems, where your company is right now in terms of capabilities and what needs to be done before you can properly implement AI to address the problems you want to solve or the opportunities you want to take.

There is no one magic algorithm that can solve all your problems: Before implementing AI, you should focus on the problems you have and how AI can help you solve them. Next, you need to check whether the data you need to solve your problems is available. Often a hybrid between algorithms can be the right fit, depending on the problems you want to solve.

Supervised learning provides the most economic value: Supervised learning is currently the most applied form of learning and provides the most value for a wide variety of applications. When beginning to implement AI in your enterprise, you'll likely be working with supervised learning.

In the next section of this paper, we'll talk about how IBM can play a role in implementing AI within your organization. IBM not only has many years of experience with these types of projects but has also been a pioneer in the AI arena. Based on its experience and knowledge, IBM can help companies of all sizes implement AI solutions.

The IBM Garage as the engine for a Cognitive Enterprise transformation

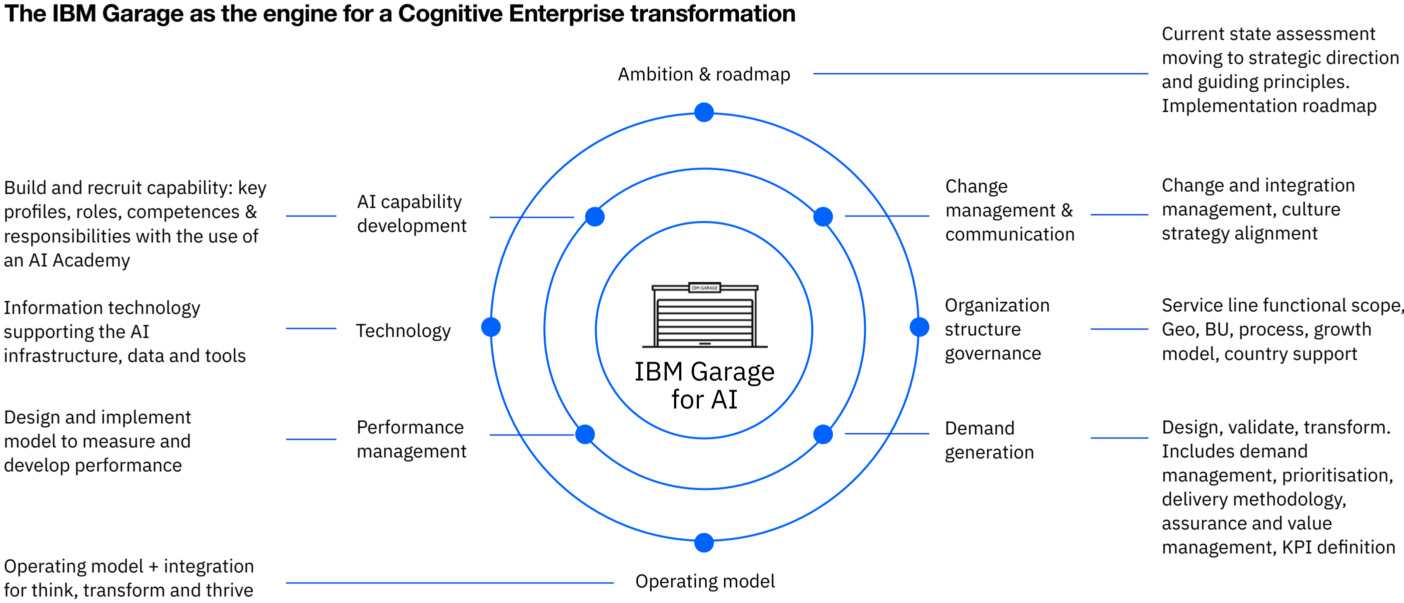

Figure 6: The IBM Garage for AI.

IBM Services

IBM has the experience and knowledge to help guide your company through a business and technology transformation. IBM puts this into practice through the “garage” concept. The IBM Garage lets you experiment with big ideas, acquire new expertise and build new enterprise-grade solutions with modern and emerging technologies for immediate market impact.

IBM looks at implementing AI in a holistic way. You as a client and IBM can enter into a strategic alliance to transform your business by creating a platform for continuous innovation. The IBM Garage lets you innovate and develop with the speed of a start-up, at the scale and rigor of an enterprise. It offers an innovation space where clients and IBM work side-by-side to create first-of-a-kind strategies and solutions. You can then develop the expertise to transform the way you work.

Plan on a page to start the journey with the IBM Garage for AI

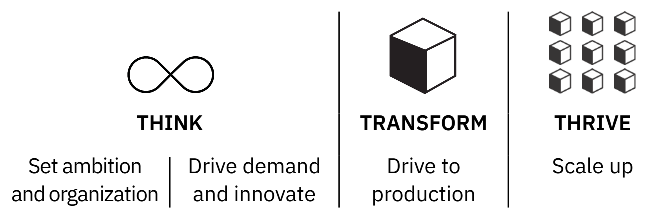

Figure 7: The IBM stages of implementing AI.

The "Think" phase is a process dedicated to thinking, experimenting, and proving, with the end user at the heart of your innovation. We start with ideas and get to working concepts fast, incorporating feedback in real time.

IBM helps you get prepared first, in a fast, pressure-cooker environment. This helps you to kickstart your AI journey. IBM works with you to deliver an assessment of your current AI capabilities.

In assessing your current AI capabilities, IBM looks at many different aspects, some of which are:

- Systems: What does your innovation system look like?
- People: How skilled is your staff in AI?
- Organization: How do you position AI expertise in your organization?
- Culture: Do you operate in an agile fashion?
- Data: How readily available is your data?
- Technology: Does your current architecture enable AI?

IBM works with you to define your AI Ambition as a high-level AI strategic direction by looking at "big ideas" and by reflecting on persona's and their user stories. A persona can be a marketing manager or a product developer. This ambition is translated into a roadmap and a high-level business case. IBM also supports you to select the most favorable operating model to move forward.

Once you are prepared for your garage, IBM uses design thinking to identify and define use cases. All use cases are conceived based on the notion of creating business value. In IBM's experience, design thinking is the most effective way. This so-called "demand generation" for the garage should be on-going, creating the backlog and setting the capacity planning for the garage.

Once the first use cases have been agreed, the garage provides the playground for deep technology, accommodating enterprise-scale ways of working—such as design thinking, agile, DevOps and lean IT. Innovation starts with a Proof of Value (POV): evaluating the business and technical feasibility of the use case in a time-boxed agile manner. Agile development is at the core of what IBM does in the garage. The goal is to succeed or fail quickly, with speed to scale. Flexibility and speed are key! A successful POV will move into a pilot implementation. The IBM Garage leverages IBM's extensive asset library to more efficiently validate use cases and accelerate transformative change.

Organizations should be prepared to step out of their comfort zones and think differently. Thinking in a data-driven and AI approach requires a shift in the way we see. It takes some effort to change our ways of thinking. This is where digital change management plays a role. You also need to consider the buildup of AI capabilities in your organization. IBM offers various forms of training and enablement ranging from Hackathons, planned learning universities for upskilling and an AI learning academy.

In the "Transform" phase, IBM collaborates with experts, data and emerging technologies, using accelerators to build minimum, viable products into production and realize business outcomes and customer adoption within weeks.

The third and last phase is the factory "Thrive" phase, which is designed to rapidly scale solutions while establishing methods and new ways of working across your enterprise that can last a lifetime. The thrive phase works as a managed service to embed and maintain data and AI solutions at scale. Using a factory-like approach, you put into production something you want to be consistent in design and quality, and that doesn't require innovation in execution. IBM focuses on successful delivery of projects aligned to the roadmap and AI platform technology and helps ensure a clearly defined governance.

Using this method, IBM has been successful many times and wants to keep improving the method with each project delivered. IBM is one of the frontrunners of AI and AI implementation, and thrives on applying its knowledge and experience to help improve the world.

About the authors

Jorn Jansen Schoonhoven, a data scientist at IBM with two years of data science experience is part of the Advanced Analytics branch of the IBM Amsterdam office, and part of the IBM Global Institute of Business Value team. He holds a master's degree in business analytics and big data, and a master's degree in management from IE Business School. He is the main author of this paper and can be reached at Jorn.Jansen.Schoonhoven@ibm.com or +316 22403033.

Marloes Roelands is an Associate Partner with over 20 years' experience in consulting. She is the European leader for the "IBM Garage for AI." She loves making innovation happen with her clients and followed the Executive programme "Strategy and Innovation" at Saïd Business School at Oxford University to support her thinking. She also holds a Master of Economics from Erasmus University Rotterdam. She can be reached at marloes.roelands@nl.ibm.com.

Francesco Brenna, an Executive Partner with over 17 years of consulting experience, currently leads the AI practice for IBM Global Business Services in Europe. He holds a Bachelor of Science in Computer Science from Zurich University of Applied Sciences in Zurich and a Master of Business Administration (with distinction) from Warwick Business School. He can be reached at francesco.brenna@ch.ibm.com.

Acknowledgements

The authors thank the advanced analytics practice for sharing their experience about the implementation of AI, the IBM Institute for Business value for their advice and insights, and other IBM branches whose insights are reflected in this document.

The authors are also grateful for the contributions of many IBM colleagues. In particular, they thank Mando Rotman, Wouter Oosterbosch, Damian Brennan and Vanessa van de Vliet. This report would not have been possible without the help of many other colleagues within the IBM Watson and advanced analytics department.

Finally, special thanks go to Apostolos Mourouzis and Patricio Fernandez Weisson.

Appendix

Different types of learning

Learning. which is one of the fundamentals of artificial intelligence and machine learning, is when the algorithm improves itself by looking at the data provided. There are two elements involved: knowledge and feedback. Knowledge provides information that's already in the data, and the algorithm can learn from feedback through interactions with the user. This happens when a user gives the model feedback about correct or falsely predicted outcomes. There are four types of machine learning: supervised, unsupervised, reinforcement and transfer. Currently the most often used type is supervised learning, and thus we can say that the most economical value is created within this category.

Supervised learning is a learning method that maps an input to an output using human data and feedback to improve. A data set is provided with associated correct labels to the data. An example would be pictures of animals in which all pictures were correctly labelled as the animal in the pictures. Supervised learning trains based on historical data and builds rules that can be applied to predict future problems. The better the data set, the better the output.

You may use this type of learning when you want to classify or predict outcomes. With regression, you are predicting a continuous value (“How much will the stock price be?”). With classifying, you are assigning a label to an input (“Is this picture a man or a woman?”). Other examples would be using speech recognition to examine the sentiments of people calling your customer service center, or image recognition to define products in a warehouse so they could be properly sorted.

Unsupervised learning occurs when the algorithm is not given a specific “wrong” or “right” outcome. Instead, the algorithm is given unlabeled data. Unsupervised learning is often used when you want to classify data but don’t know how to do so. For example, you’d likely use unsupervised learning if you had a set of customer data and you didn’t know what kind of classes they would fit in. An unsupervised learning algorithm can find natural groupings of similar customers in a database and the user can then describe and label them.

Reinforcement learning is a class in and of itself; it is not given a specific goal, but rather learns from trial and error. The main concept is that instead of a specific action being labelled, there is a sequence of actions that is associated with a reward. If we take a maze as an example, the algorithm will be rewarded when it comes closer to its goal and be penalized every time it gets stuck or moves away from the completion. A recent example of reinforcement learning is AlphaGo, where Google trained a deep reinforcement learning network with many examples of the game Go, eventually making its performance superior to that of even the best human. This trick is not new, since it was used in TD-Gammon in 1992, created by Gerald Tesauro at IBM. TD-Gammon was a backgammon-playing program that reached the performance of the best human players at the time.

Reinforcement learning is not currently widely used, but it does have high potential when developed more extensively. You would need a lot of data (which is not always the case and takes time to process) to be able to make reinforcement learning work.

Transfer learning is when your algorithm learns to solve one problem, takes information from this problem and then solves a new problem with that information. This currently happens a lot with image recognition. Pre-trained neural networks are used to solve new problems.

IBM Corporation
New Orchard Road
Armonk, NY 10504

Produced in the United States of America
October 2018

1 www.weforum.org/agenda/2016/01/the-fourth-industrial-revolution-what-it-means-and-how-to-respond

2 www.quora.com/What-is-artificial-intelligence-What-are-task-domains-in-AI

3 ai100.stanford.edu/sites/default/files/ai_100_report_0906fnlc_single.pdf

4 www.investopedia.com/terms/a/artificial-intelligence-ai.asp

5 neuralnetworksanddeeplearning.com/chap5.html

6 datasciencedegree.wisconsin.edu/data-science/what-do-data-scientists-do

7 www.interaction-design.org/literature/book/the-social-design-of-technical-systems-building-technologies-for-communities/the-evolution-of-computing

8 www.forbes.com/sites/blakemorgan/2017/06/13/ethics-and-artificial-intelligence-with-ibm-watsons-rob-high/#72b4a0e3260e
www.forbes.com/sites/gilpress/2016/12/21/artificial-intelligence-pioneers-peter-norvig-google/#6ecd8a2d38c6

9 www-01.ibm.com/common/ssi/cgi-bin/ssialias?htmlfid=KUM12390USEN&, page 67

10 homes.cs.washington.edu/~marcotcr/blog/lime

39019539-USEN-00