

WHITE PAPER

Smart success

Ignite the power of AI and data science across the organization



Contents

The importance of artificial intelligence	4
The benefits of introducing AI	5
Key stages within the data science project lifecycle	6
Creating a high-performing data science team	8
The critical role of internal and external partnerships	9
Five AI use cases your company can consider: Powered by OpenText Magellan	10
Realizing the benefits of using Magellan for AI	11
Resources	12

Summary

Artificial intelligence (AI) transforms organizations by empowering them to make more data-driven decisions and eliminates tedious and repetitive processes. Yet, many companies struggle to see how to deploy the right AI solution for their organization. While they may have a general understanding of what AI can do, it remains unclear as to how they can leverage data science and advanced analytics to improve efficiency.

This white paper will help you identify whether AI is the right fit for your organization and how to set up an AI structure that adds value to the business. It outlines the benefits of AI for the enterprise, provides a list of questions to ask before getting started and offers guidance for structuring a high-performing data science team. In addition, it details the iterative process for managing a successful AI project with a comprehensive, step-by-step guide.

It also provides five use cases to allow organizations to explore the potential of AI within their own organizations.



“90% of U.S. executives believe cognitive technologies are an important aspect of internal business processes.”¹

“AI is about teaching computers to do what comes naturally to humans.”

The importance of artificial intelligence

AI is transforming the way we do business across all industries. In fact, more than 90 percent of United States executives believe cognitive technologies are an important aspect of internal business processes.² Yet, many business leaders remain unsure about how to create and deploy a successful AI practice in their organizations and, more importantly, whether there is a practical application for which AI-powered solutions can add value. Business leaders know AI streamlines workflows, alleviates tedious and repetitive processes and empowers data-driven decisions, but what does that mean for your organization? How can AI help and where do you begin?

Determining whether your company could benefit from AI lies simply in asking what value you can create with AI-powered solutions. Once you answer that question, you can determine what data you will need to power those applications, which data assets you have in place and which you will need to acquire. Of course, these are very broad questions and this is where many businesses get stuck. To find answers, it may be helpful to establish a more solid understanding of what AI can do.

Simply put, AI is about teaching computers to do what comes naturally to humans. While humans can see patterns in small amounts of data, AI examines massive pools of data to quickly and efficiently extrapolate meaning and predict outcomes. With the machine learning (ML) in AI, the more data available, the more it will “learn.” AI is different from traditional software in that it has a broader range of decision-making capabilities, which makes it valuable for virtually any industry. From helping to predict when an oil rig will need repairs to analyzing customer sentiment about a new product line across various social channels, there are countless uses and benefits.

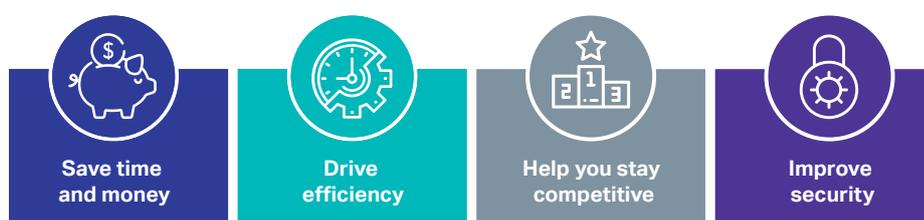
The areas in which AI and machine learning are being used most effectively are where they are leveraged not to replace humans, but instead, to help humans fulfil their roles more efficiently. Specifically, machine learning applications excel in processing and analyzing mountains of data far more rapidly than any human ever could, then returning insights with clear recommendations for human users to consider. When implemented effectively, machine learning allows users to explore possibilities and consequences of various business decisions, providing a competitive edge.

The benefits of introducing AI

With evolving customer demands and a growing sense of immediacy affecting all industries, business leaders must make decisions quicker than ever. Of course, informed decisions are born from the ability to analyze as much data as possible, but collecting, contextualizing, analyzing and deriving insights from huge volumes of information is too time-consuming.

Machine learning automates the tedious elements of human research and review to provide valuable analyses and recommendations, amplifying the role of human employees across business processes—not eliminating it, as some might fear.

With these capabilities, AI-powered solutions can:



Maximizing the benefits for your enterprise

The AI journey will be different for each organization based on its assessments of needs, strengths, processes, goals, etc. It should therefore not be expected that all organizations adopt AI in the same way. Consider the following questions when determining whether an AI-powered solution will not only be feasible, but also beneficial, for your organization:

What specific objectives can we achieve by introducing AI?

AI will only yield organization-wide success if you are able to look at it through a strategic business lens—not as an IT project. This requires cross-functional collaboration to identify the organization's overarching objectives. You might choose to implement technology in a way that reduces risk, improves productivity or helps you make more accurate predictions. While AI can do all of this and more, starting with a specific goal will help you develop a detailed and actionable plan to get there.

How will AI help us secure a competitive edge?

If you are not sure how to leverage AI to help secure better performance, consider any industry trends that might be taking shape. Analyzing the direction your business line is moving in should help you identify ways to meet customer demands and stay competitive with AI. In Financial Services, for example, banks are using AI to help them make more tailored product and service recommendations. Use cases from your sector will also prove to be a valuable resource to help you envision the right solutions for your needs.

Do you have the right data?

Do you have data relevant to answering the business problem? How much data do you have? Where is your data? Is data readily accessible for analysis and modelling?

With the answers to these questions in mind, you can begin pursuing a successful AI project in your organization.

1. Define task
2. Collect data
3. Model exploration
4. Model refinement
5. Testing and evaluation
6. Deployment and integration
7. Continuous maintenance

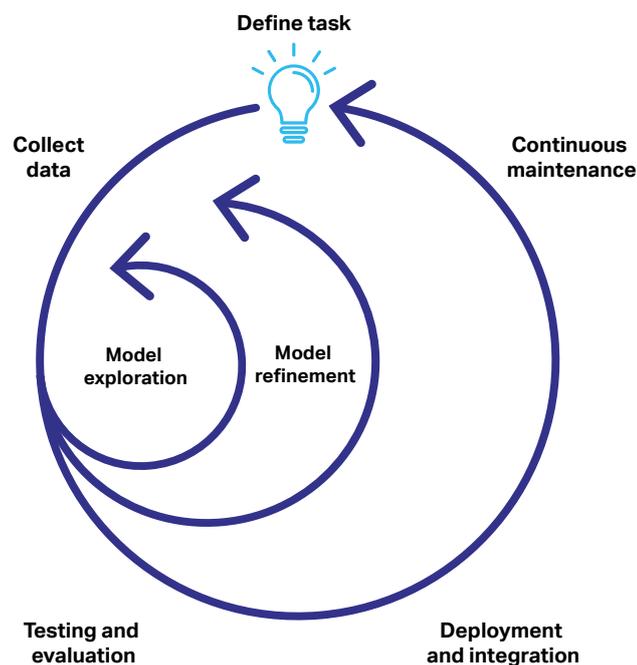
Key stages within the data science project lifecycle

The process of developing a machine learning model is highly iterative. Often, you will find yourself returning to previous steps before proceeding to a subsequent one. After all, a machine learning project is not considered complete after the first version has been deployed. Instead, the feedback you collect after the initial version will help you shape new goals and improvements for the next iteration.

1. Project planning/Define task

Many AI projects fall short because organizations jump directly into modeling without taking the time to plan. Starting with a cohesive plan gets everyone on the same page and may prevent roadblocks down the road. While it is not always simple to define a model task from the beginning, brainstorming should help you identify projects with high potential, including those that could make a big impact, yet still be feasible. Where might prediction drive the largest value?

During this step, you should also establish what kind of data you will need to support your AI project, how you will acquire it and how much data is needed. For example, if you are thinking of deploying a solution for predictive maintenance on a train, you might pull information from IoT sensors, weather patterns and passenger travel patterns.



Machine learning development lifecycle

2. Data acquisition/Collect data

How you label your data can have a huge impact on model performance. In some cases, you may be able to use a self-labeling system, but in others, manually labeling data may be necessary. Be consistent in your labeling criteria and be sure to associate specific models with dataset versions to avoid confusion.

3. Model exploration

To determine whether the models your data scientists are creating will be effective, you must first determine baselines against which model performance will be measured. The best approach is to start small with an initial data pipeline. You will likely test many ideas throughout this step, but ultimately, one leading model will emerge. You can reproduce the results of this model, then apply it to your dataset for a second baseline. Thereafter, you can revisit steps one and two to ensure it fits.

4. Model refinement

Once the model is in place, refine it by debugging it and performing an error analysis to pinpoint potential failures. As you enhance the model, consider the following questions:

- Can the data scientist explain their choices and are they reproducible?
- How is the accuracy measured?
- Have you measured the potential profit savings of the project?
- Can others understand the notes and project methodology?
- What is the release plan and how will you measure success?
- Has data security been reviewed on the algorithms and pipelines?

5. Testing and evaluation

Often, data scientists find that “data in the wild” tends to be different from the data they used for training purposes. For this reason, it is important to revisit the metric for model evaluation to ensure that the one you have chosen is appropriate for facilitating the desired actions. Now is also the time to write tests for the input data pipeline, model inference functionality and specific scenarios expected in production.

6. Deployment and integration

For best results, start with a model deployment among a small group of users. If everything runs smoothly, roll it out to all users. You should continue to monitor live data and model prediction distributions.

7. Continuous maintenance

Following deployment, be sure to address any changes that impact the system and retrain the model periodically to prevent it from becoming worn out. Should model ownership transfer, you should ensure that the new team is properly educated on the model.

Of course, to execute these steps effectively, you must have the proper team in place.



Creating a high-performing data science team

AI projects are not solo sports. They require the ongoing contributions and collaboration from a number of key team members. Because data science is an intersection of various functions, including math and statistics, computer science and business knowledge, you will need a dynamic group of individuals to form a high-performing machine learning team.

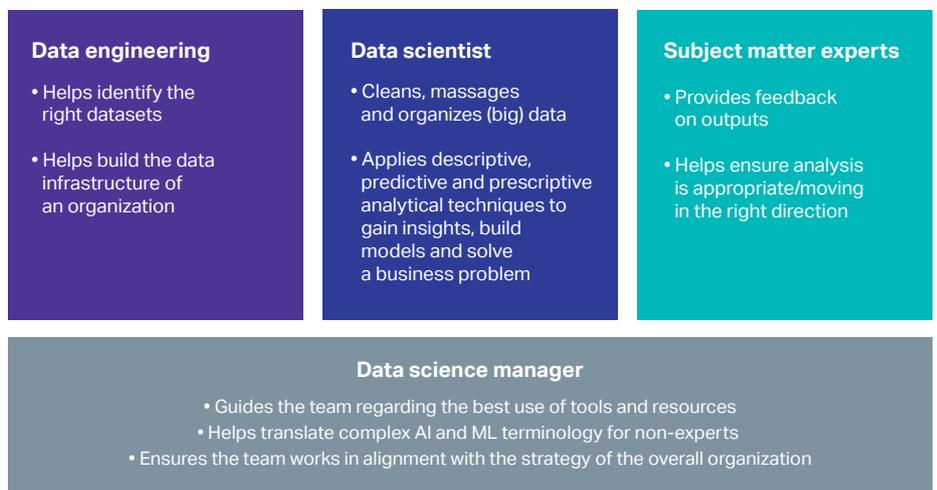
Typically, teams involve individuals from data science and data engineering, as well as subject matter experts. Here is a look at each of these roles in more detail:

Data engineers: These specialists are typically tasked with building the data infrastructure of an organization. They have strong programming and hardware skills, are familiar with big data technologies and excel in building data pipelines. They are not necessarily experts in analyzing and modelling data but should be able to work with relevant business divisions to determine the data characteristics for each use case.

Data scientists: After you have built out your data infrastructure, you will need people who can take that data, clean it, analyze it, apply algorithms, run experiments and communicate results effectively. These professionals typically use tools, such as Jupyter™ notebook and RStudio®, have knowledge of programming languages, such as R and Python®, and experience working with big data technologies, such as Apache Hadoop® and Apache Spark™. In addition, they have a strong background in statistics, programming and machine learning.

Subject matter experts: Subject matter experts help cultivate strategies, generate ideas and examine factors necessary for supporting different use cases. They understand the key business problem, processes and challenges. Their input and knowledge are crucial in helping the team develop meaningful solutions that provide business value and overcome challenges.

Data science manager: Data science managers are hands-on leaders who help build the foundation of an organization's data science strategy, recruit and create talent teams, ensure effective communication among members and develop processes for the team to follow. They are in charge of connecting the data science and analytics team with the rest of the organization, other divisions and executives. One of their core responsibilities is to translate complex AI and ML terminology to non-experts and make sure the team works in alignment with the strategy of the overall organization.



Data science core team



The critical role of internal and external partnerships

For effective execution, AI must be built on strong partnerships, both internally and externally. Each party involved must be aligned with the overarching data strategy to keep all efforts focused on the same goal. When everyone is aligned and knows their role, they can achieve real, measurable results from the AI practice.

Internal partnerships

Not only is having the right people on your team important, so is ensuring they have the resources needed to adopt AI technology and to collaborate. To facilitate broad AI training, make sure you have input from an AI expert who can guide your organization. They can work with your chief learning officer (CLO) to find or develop training materials, including video tutorials and courses. The organization should then develop procedures for each role involved with AI, including executive leaders, business unit managers and data scientists.

Before, during and after deployment, there must be constant communication between the AI team and the business unit. Each specialist has strengths and boundaries to their skillsets, so working with others will be an essential aspect of project success. Typically, data scientists or data science managers will ensure cross-functional collaboration to create business value and ensure complex AI terms are being translated to usable insights. These individuals must also ensure that their work is aligned with the organization's overarching goal. This is where AI systems empower teams to break down silos and communicate insights across teams.

External partnerships

Finding the right AI partner for your organization is another component integral to the success of your AI practice. A trusted AI partner should allow you to leverage underused yet valuable data, allow data scientists to work in their familiar environments, such as Jupyter notebook and R, and provide the ability to operationalize models in business-friendly interfaces. Establishing a project manager who can collaborate with your AI partner will also be critical to success.

The importance of executive buy-in

One common cause of failure among AI projects is lack of buy-in. While data scientists and subject matter experts may be able to see the value of AI in the enterprise, support must start at the top. Communicating how an AI investment could help the company save money and piloting solutions can help teams establish credibility and attract support from leadership. Moreover, explaining model outputs in concise and tangible terms can help accelerate adoption and build trust among stakeholders.

Five AI use cases your company can consider: Powered by OpenText Magellan

AI-augmented recommendations

Consumers are inundated with frequent and irrelevant product offers. To understand how they could better serve their customers, a major bank deployed a customized personalization model with intelligent recommendations. The bank implemented a system that harnesses AI and big data to monitor customer actions, predict their needs and deliver highly personalized experiences across every channel. With a deeper insight into customer preferences and behaviors, the bank has improved deals and offers from retail partners. With fewer annoying and irrelevant offers, the customer experience has improved and the bank has seen higher click-through and conversion rates on their offers.

Intelligent capture

While traditional content capture systems help businesses manage their content, there is a growing need to automate the classification and routing of content at scale. AI-enhanced capture can extract information from enterprise content and route documents to the appropriate workflow. With the solution, organizations, such as large banks, can intelligently classify documents and free up staff from sorting through thousands of incoming documents every day.

Classification/Smart migration

From spreadsheets to Microsoft® Word documents and email, organizations are swarming with billions of pieces of content. Managing this unstructured, textual data is notoriously difficult, but machine learning can help organizations categorize information and trigger processes, thereby automating knowledge classification. This is especially beneficial for organizations with sensitive data, including personally identifiable information (PII), as AI helps companies pinpoint where this information is located and can support compliant activities surrounding its handling. Some companies are also using AI to sort through “Contact Us” inquiries and feedback. Using advanced text analytics, AI systems can recognize the nature and tone of customer submissions and classify this textual data to trigger workflows accordingly. For example, complaints might receive a higher priority status than general comments.

AI-powered predictive maintenance

For industrial plants, transportation machinery and facilities with high-value assets, minimizing downtime and reducing maintenance costs are essential to staying competitive. While many companies in industries, such as logistics, vehicle and fleet management, energy and utilities and transportation, leverage traditional IT solutions for asset management, these systems do not have the AI capabilities needed to derive value from unstructured data. With AI-powered solutions, companies are able to acquire, merge, manage and analyze big data from sources, such as IoT sensors, to receive predictive insights that improve over time as assets are used and maintained. As a result, companies that leverage AI for asset performance optimization increase the efficiency of high-value equipment, reduce downtime, decrease unexpected repairs and lower maintenance costs.

Voice of the Customer

Nearly 80 percent of all enterprise data is unstructured.³ Within these emails, memos, social media streams and customer service chats, there is real business value that organizations can leverage. Yet, doing so manually is tedious and time-consuming and traditional analytics systems cannot interpret the nuances of human language. Specialized AI algorithms can assess human sentiment rapidly across a broad range of sources. With an AI-powered analytics platform, organizations can improve the process of tracking customers through the buying journey. Gathering insights from social media, customer surveys and other sources, it captures brand-related concepts and delivers relevant, valuable

Realizing the benefits of using Magellan for AI

OpenText™ Magellan™ combines machine learning and predictive analytics on one cohesive platform to help organizations extract the greatest value from their big data and big content. It allows for deep dives into structured and unstructured data with self-service insights powered by AI. Plus, with support for popular data science coding languages, such as Python, Scala and R, it enables data scientists to work in their familiar environment while sharing machine learning models across the enterprise with the Magellan data science notebook. Subject matter experts and business users can then explore information through analytics dashboards and other user-friendly data visualizations.

Here are just a few ways companies use Magellan to drive operational efficiencies:

Drive productivity

Automate repetitive, tedious tasks by reducing manual processing and augmenting enterprise data management with AI and analytics

Reduce costs

Minimize the costs of a technology investment by eliminating disparate tools with limited processing power and upgrading to open source AI software that caters to big data analytics

Enhance visibility and decision-making

Empower business users to dive into massive amounts of structured and unstructured data with advanced AI analytics functionality

Leverage convenient drag-and-drop features

Allow employees across the organization, including subject matter experts and business users, to explore data through simple yet effective, drag-and-drop advanced analytics workflows

Improve efficiency

Leverage machine learning to unlock the value within enterprise data to increase operational efficiency through AI

Achieve continuous improvement

Rely on a trusted software leader with an unrivaled track record in enterprise information management to digitally transform the enterprise with open source analytics tool

[Watch the webinar](#)

[Read the solution overview](#)

[Read the infographic](#)

[Join the conversation](#)

Blog: 7 steps to getting started with your AI journey

Blog: Tapping your stores of unstructured data can yield new insights

Initiate a successful AI practice with OpenText

OpenText is a global organization that can work with your team on its AI journey. The Magellan team includes data scientists, computational linguists and experts in machine learning, text mining and analytics and data engineering. They have deep expertise with Magellan and other advanced analytical technologies and are experienced in working with organizations across many industries to gain actionable insights from their data. They can help you take your project from a proof of concept to enterprise deployment.

As an international company, OpenText understands the global challenges of data regulations, such as the GDPR. OpenText has thoroughly analyzed GDPR requirements and has implemented an organization-wide compliance strategy to meet all requirements in the regulation and to help customers do the same.

Choosing OpenText allows you to connect with a wide selection of proven enterprise tools and services to help your organization lower operational costs, reduce risks and accelerate business. Based on a recent survey by O'Reilly, the top blockers to AI within an organization are company culture, difficulties identifying use cases and a lack of skilled people.⁴ OpenText Professional Services provides tools to overcome these barriers, such as cognitive strategy workshops and data science project methodology training.

About OpenText

OpenText, The Information Company, enables organizations to gain insight through market leading information management solutions, on-premises or in the cloud. For more information about OpenText (NASDAQ: OTEX, TSX: OTEX) visit: opentext.com.

Connect with us:

- [OpenText CEO Mark Barrenechea's blog](#)
- [Twitter](#) | [LinkedIn](#)

¹Deloitte, The 2017 Deloitte State of Cognitive Survey. (2017) <https://www2.deloitte.com/content/dam/Deloitte/us/Documents/deloitte-analytics/us-da-2017-deloitte-state-of-cognitive-survey.pdf>

²Ibid.

³Forbes, The Big (Unstructured Data Problem). (2017) <https://www.forbes.com/sites/forbestechcouncil/2017/06/05/the-big-unstructured-data-problem/#2ebcd96c493a>

⁴O'Reilly, How Companies Are Putting AI to Work Through Deep Learning. (2019) <https://www.oreilly.com/radar/overcoming-barriers-to-ai-adoption/>