451 Research® | PATHFINDER REPORT

Accelerating AI with Data Management; Accelerating Data Management with AI

COMMISSIONED BY

IBM

JULY 2019

About this paper

A Pathfinder paper navigates decision-makers through the issues surrounding a specific technology or business case, explores the business value of adoption, and recommends the range of considerations and concrete next steps in the decision-making process.

ABOUT THE AUTHORS

MATT ASLETT

RESEARCH VICE PRESIDENT

Matt Aslett is a Research Vice President with responsibility for 451 Research's Data, AI and Analytics Channel – including operational and analytic databases, Hadoop, grid/cache, stream processing, data integration, data governance, and data management, as well as data science and analytics, machine learning and AI. Matt's own primary area of focus currently includes distributed data management, data catalogs, business intelligence and analytics, data science management, and enterprise knowledge graphs.

JAMES CURTIS

SENIOR ANALYST, DATA, AI & ANALYTICS

James Curtis is a Senior Analyst for the Data, AI & Analytics Channel at 451 Research. He has had experience covering the BI reporting and analytics sector and currently covers Hadoop, NoSQL and related analytic and operational database technologies.

Executive Summary

Artificial intelligence (AI), including machine learning and deep learning, is set to become one of the most transformational technologies in the history of the world, affecting most aspects of our lives whether we're conscious of it or not. The application of these technologies will likely reshape how people work, study, travel, govern, consume and pursue leisure activities.

Successful AI relies on a number of factors including a large corpus of data, the requisite algorithms, expert data scientists with appropriate skills, and appropriate compute resources. The latter means not just physical and virtual server infrastructure, but also data management and database software designed to support high-performance data processing and analytics.

Data management is a critical enabler of machine learning projects because it helps overcome challenges such as accessing and preparing data, which can be a significant barrier to success. The results of 451 Research's Voice of the Enterprise: AI & Machine Learning survey, conducted with people directly involved in AI and ML initiatives, illustrates the point: 33% of respondents cited accessing and preparing data as a barrier the use of machine learning, and 15% cited it as the most significant barrier.

Figure 1: Barriers to using machine learning

Source: 451 Research's Voice of the Enterprise: AI & Machine Learning 2H 2018

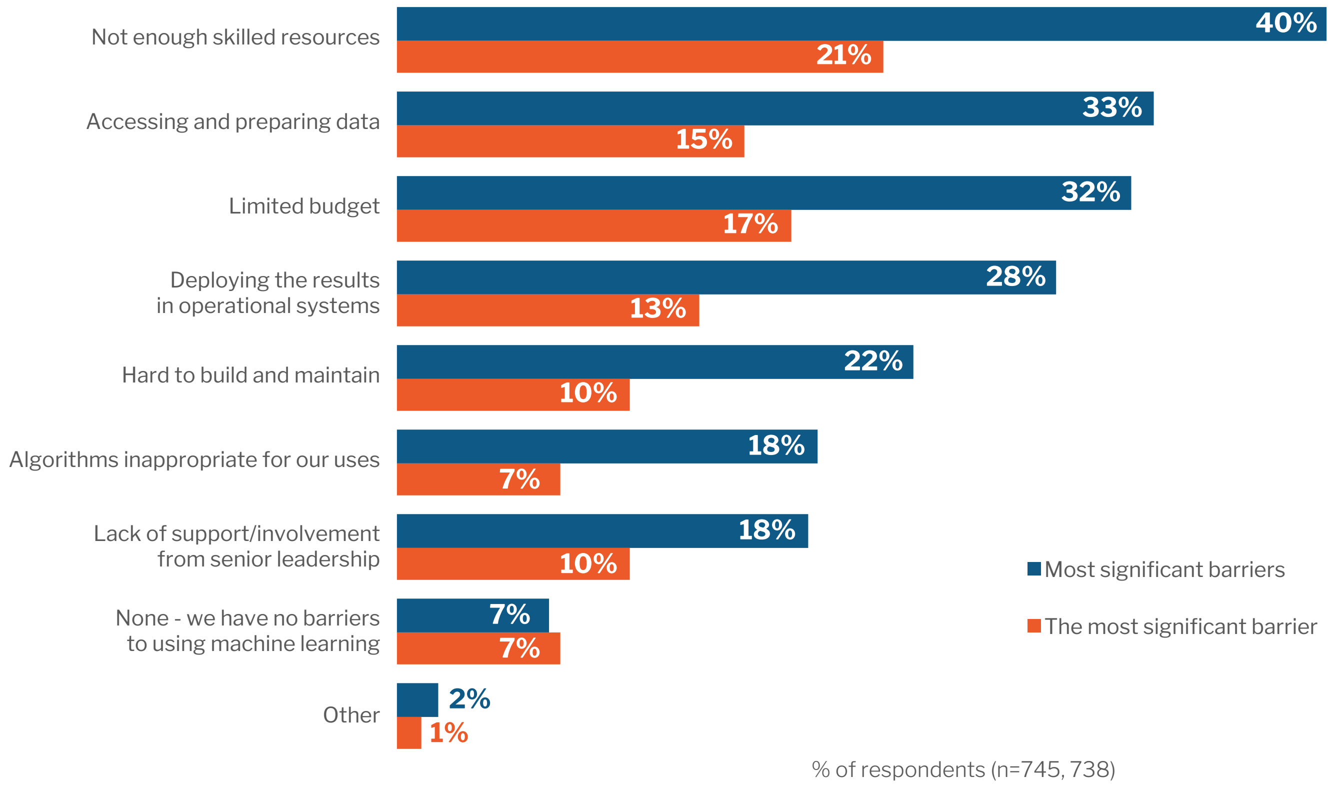

While data management can improve the development of AI applications, AI can be used to improve data management in areas such as data ingestion and query performance, enabling data engineers to accelerate data management and analytics projects and database administrators (DBAs) to focus on higher-impact tasks.

In order to take full advantage of everything that AI has to offer, enterprises must embed AI at multiple levels, starting at the data level to ensure that AI enables the full scope of the data management lifecycle, from ingestion to curation and discovery, as well as driving applications that are built to leverage that data.

The results of 451 Research's Voice of the Enterprise: Data Platforms and Analytics survey, conducted with people responsible for data platforms and analytics initiatives, reveals the extent to which enterprises see AI and ML as critical aspects of their data and analytics projects. Two-thirds of all respondents agree that AI and ML are an important component of their data platform and analytics initiatives, but this figure increases to 88% among the most data-driven companies (i.e., those at which nearly all strategic decisions are data-driven).

Figure 2: Importance of AI and machine learning to data platforms and analytics initiatives

Source: 451 Research's Voice of the Enterprise: Data and Analytics, 1H 2019

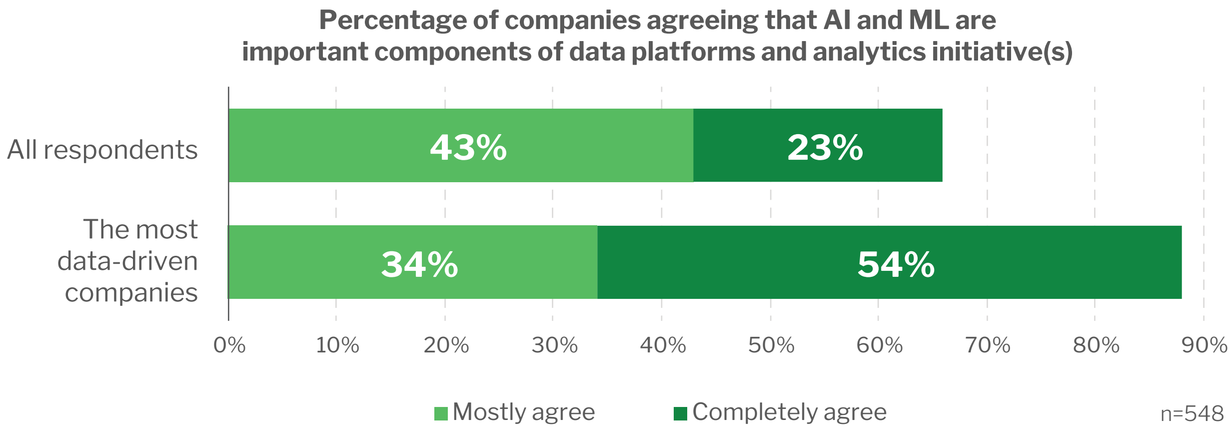

In relation to data and the systems that manage it, data architects and DBAs are challenged on multiple fronts: they must increase operational efficiencies while providing data access to a greater variety of data consumers – including data and business analysts, senior executives, developers and data scientists. In order to do so, database administrators need data management systems that run efficiently at high performance, capable of producing accurate results.

Additionally, enterprises need the data to be easily accessible to data scientists for building AI-enabled applications. As the underlying data platforms evolve to better support AI initiatives, these systems need to provide support for multiple languages such Python, GO, JSON and Jupyter notebooks so that the development of AI-based applications and the construction of complex data models can be accelerated.

The Synergistic Nature of Data Management and AI

Data management systems and AI are synergistic. When AI becomes embedded throughout the data management system, it has the potential to improve database query accuracy and performance, as well as optimize system resources, reducing the burden on DBAs while improving data access for data scientists and developers.

Enterprises face increasing challenges as they look for ways to improve the operational efficiency of their data management systems while enabling greater data access to a variety of data consumers, particularly data analysts and data scientists. One approach is to adopt point products that have AI capabilities built in. However, that may address only one aspect of the data management environment. It may also add overhead in terms of copying and moving data from one product to another. Another approach is to embed AI throughout the environment, specifically at the data layer, that serves as the common foundation for all data consumers.

When AI gets embedded at the data layer, it creates a synergistic relationship between the underlying data management system and the development of AI applications, which has the potential to impact the entire data lifecycle. That is, database administrators and architects can better manage and oversee the data, which then leads to high-quality data that can be accessed more efficiently by data scientists and developers who are building applications and services.

Data from 451 Research's Voice of the Enterprise: AI & Machine Learning survey points to the increasing adoption of machine learning by enterprises. More than two-thirds of enterprises have already adopted machine learning or have plans to do so within the next two to three years. Perhaps more significantly, of the 40% of respondents who said they already have machine learning in use or proof of concept, 92% have positive opinions about the performance of their machine learning projects, with 37% rating them as very successful and 55% rating them as somewhat successful.

Figure 3: Enterprise adoption of, and attitudes toward, machine learning

Source: 451 Research's Voice of the Enterprise: AI & Machine Learning 2H 2018

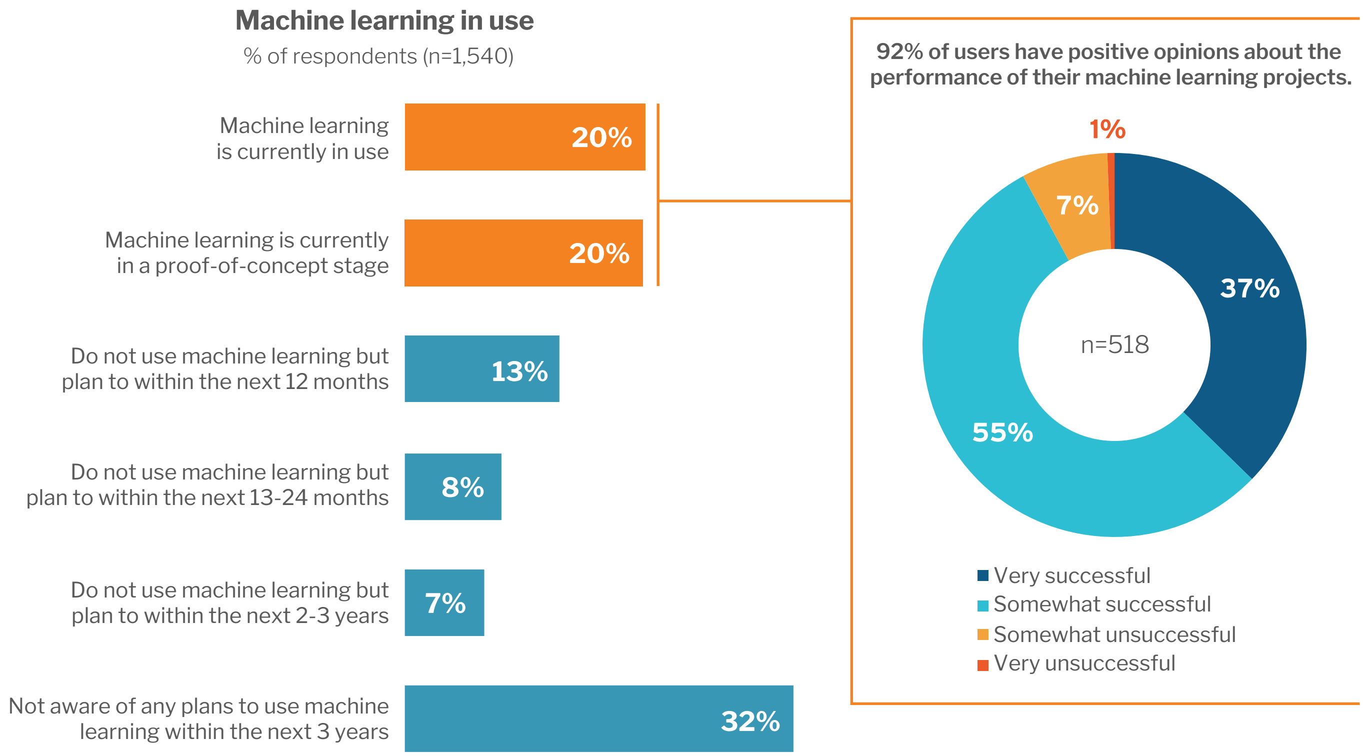

There are many use cases for artificial intelligence in multiple industry sectors. At a high level, those use cases can be considered part of a spectrum: At one end of that spectrum, AI is deployed to automate highly complex tasks, and at the other end of the spectrum, AI is used to automate predictable and repetitive tasks. In the middle, it can be used to augment humans doing everything in between.

Figure 4: The spectrum of AI use cases

Source: 451 Research

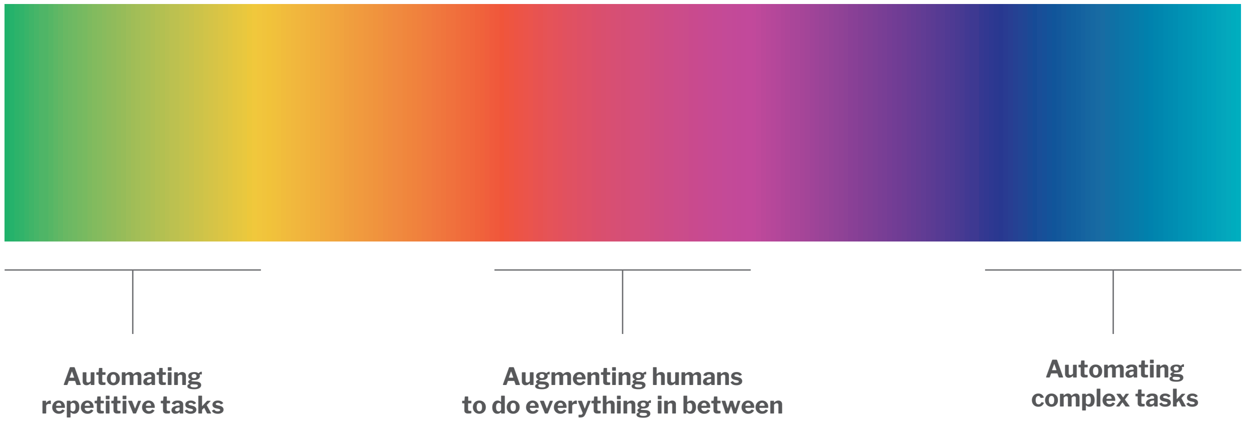

This spectrum can be used as a lens through which to examine the multiple ways AI can be applied to data management, and vice versa. While this is a high-level view of the ways in which AI and data management are synergistic, there are multiple examples of each, examined in detail below.

Figure 5: The spectrum of AI use cases as they relate data management

Source: 451 Research

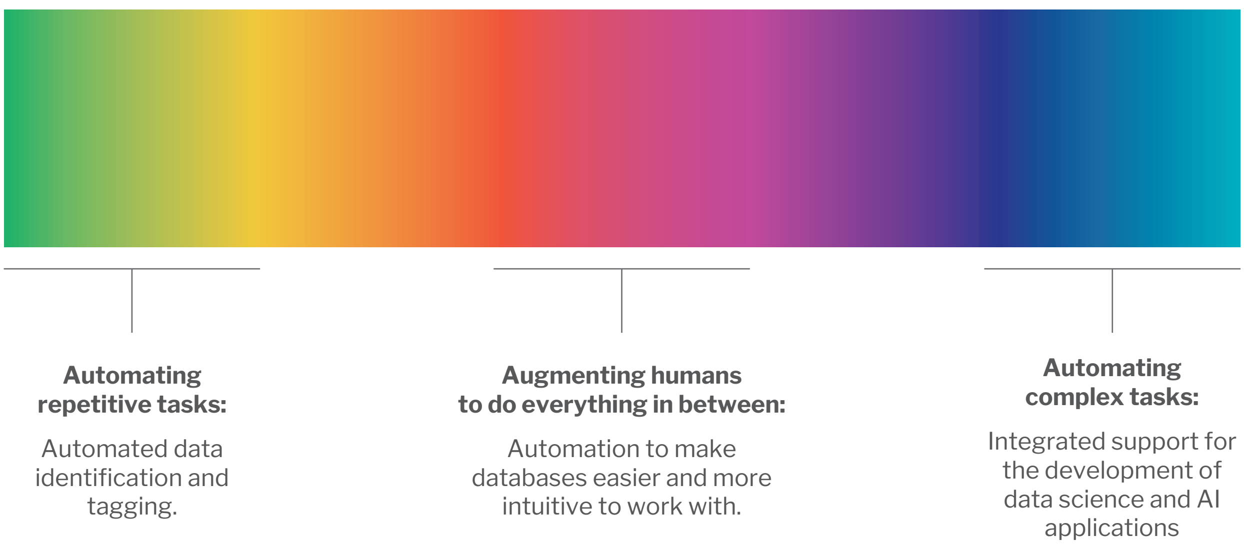

Automating Repetitive Tasks

As noted above, data ingestion and preparation are fundamental aspects of the AI data pipeline. They are also some of the most laborious and time-consuming tasks, making them prime candidates for automation. Machine learning has a key role to play in improving the efficiency of this stage of the pipeline by automating the identification and tagging of data to reduce the need for manual data preparation.

While databases and good data management are important factors for the acceleration of AI, there are also potential advantages of using AI to improve data management, particularly through the automation of tasks that are repetitive and somewhat predictable. One area of focus is improving the efficiency of databases – query optimization, in particular. There are two aspects to enhance query processing: performance and accuracy. Neither can be considered mundane – in fact, they both require a high degree of sophistication from data engineers and DBAs – but they are highly repetitive and manual tasks that have the potential to be accelerated through automation.

Regarding performance, enterprises often struggle to ensure that database systems are running efficiently. Queries that overload the system, consume excessive resources or impact other running jobs not only affect performance but also require manual resources to rectify. AI can help by automating the management of these queries based on resource consumption, providing a more stable and reliable system that can prioritize queries, reducing manual governance and monitoring of the database.

In relation to query optimization, an AI-enabled database can have a dramatic impact on improving the accuracy of estimating cardinality related to planned queries and creating a query plan to deliver the greatest efficiency. Applying machine learning to the query-optimization process can result in ongoing cardinality estimation improvement driven by constant feedback. The overall impetus is that by executing queries in a more optimized manner based on AI methods, enterprises not only can decrease the time it takes to generate insight but also make more confidence-based business decisions.

There are other areas in which AI can automate common and predictable tasks, and these areas can overlap with what is considered human-augmented tasks. The difference, however, lies in the level of human involvement. In relation to query optimization (described above), there is little to no human involvement. Other focus areas where reducing human intervention might be considered beneficial include applying security updates to guard against external attacks, managing database uptime and availability and anticipating infrastructure failures, and database management such as monitoring and tuning to drive performance.

Once again, the aim is not to eradicate human involvement, but to take advantage of automation to enable database administrators and data engineers to focus on higher-impact tasks.

Augmentation of Everything in Between

There are two key areas where AI at the database layer can augment humans to improve the overall operational efficiency of the business and accelerate business decision-making: empowering business analysts and evolving the role of the database administrator. Regarding the former, one of the primary challenges related to analytics projects has historically been to 'democratize' business intelligence by enabling a broader range of people to make analytics-driven decisions.

To put this in context, 451 Research estimates that there are currently between one and two million data scientists in the world, compared to roughly 5-10 million business intelligence 'power users,' 50-65 million business intelligence 'casual users' and 200-250 million knowledge workers.

Figure 6: Potential number of data scientists and business intelligence users/consumers

Source: 451 Research

While data scientists serve an undeniably valuable role, their relative paucity means that they are expensive to hire and retain. There is arguably more to be gained, in terms of the depth and breadth of data-driven decision-making required to deliver effective results, from ensuring that the intelligence resulting from data science projects makes its way into the hands of business intelligence users and knowledge workers.

Databases and data management software designed to complement data science compose the engine that can accelerate the development of AI-based applications, supporting the development of a greater number of applications and enabling the output of machine learning models to more rapidly be made available to domain experts and business decision-makers. For example, the ability to call machine learning models from SQL as user-defined functions can help to operationalize machine learning projects, making the use of machine learning models transparent to data engineers and business analysts.

The automation and acceleration of data processing also has implications for data engineers and DBAs. The automation of database management is an ongoing process, but repetitive tasks that can already be at least partially automated include database provisioning and patching; performance tuning; backups, high availability and disaster recovery; database optimization; query optimization; and schema changes.

This has implications for the role played by the data engineer/DBA, although redundancy is some way off yet. We see parallels here to the development of driverless cars. Although there is a lot of talk about autonomous vehicles, for mainstream use cases today (and for the foreseeable future), the reality is closer to driver assistance and partial or conditional automation.

In the same way, automation enables the augmentation of DBAs rather than their imminent replacement. Specifically, there are several higher-impact tasks that cannot currently be automated, such as architecture planning and data modeling; data security and lifecycle management; application-specific tuning; service level management; test data management; and database sharding.

Automating Complex Tasks

As noted above, 451 Research's survey results indicate that accessing and preparing data is one of the most significant barriers to ML adoption. Furthermore, it is clear that data management has a fundamental role to play in the development and deployment of AI-based applications. As the illustration of the AI data pipeline (see Figure 7) shows, the development of AI applications begins with the ingestion and preparation of data. Database and data management products that accelerate data exploration can, therefore, lay the foundation for more efficient development of AI applications.

Figure 7: The AI data pipeline

Source: 451 Research

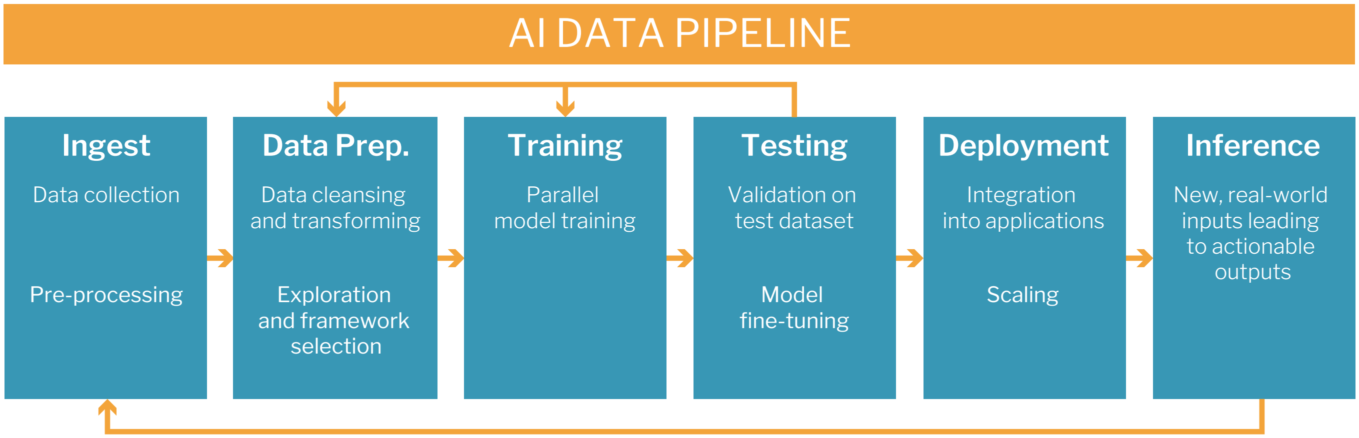

The significance of the data ingestion and preparation stage cannot be overlooked since 39% of respondents to 451 Research's Voice of the Enterprise: AI and Machine Learning survey said they believe this stage of the AI process to be the most demanding in relation to their underlying infrastructure, compared to 34% for training and 27% for inferencing. As such, by selecting and configuring a database that is designed to support more efficient data ingestion and preparation, data architects can help accelerate the development of AI applications.

Figure 8: Most demanding stage of AI process

Source: 451 Research's Voice of the Enterprise: AI & Machine Learning 2H 2018

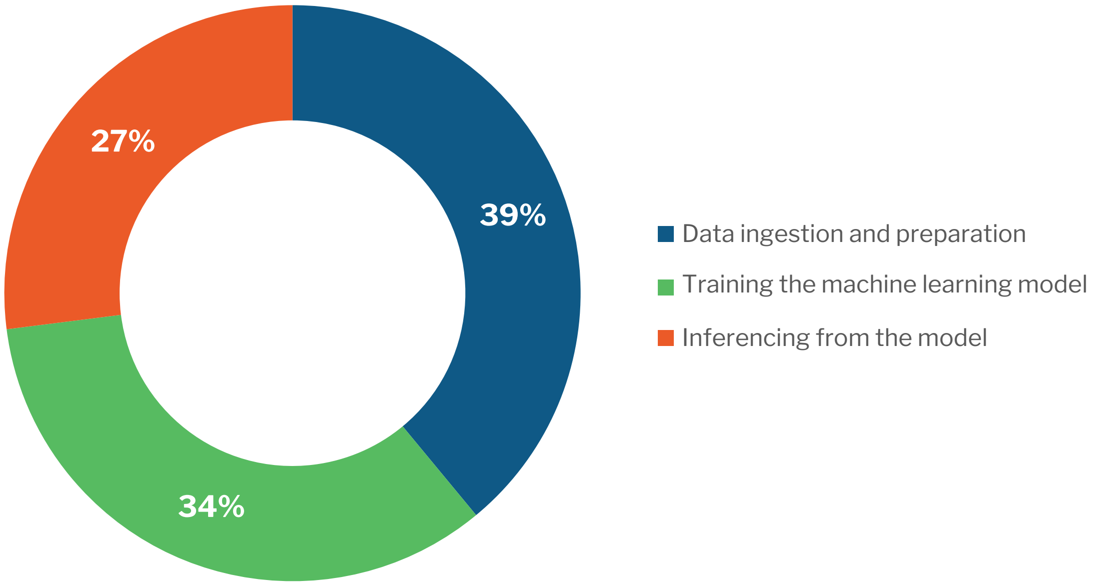

As the AI data pipeline (Figure 7) illustrates, however, data ingestion and data preparation are not restricted to the start of the development process, but continue through the training, testing and inference stages. Multi-modal databases that support a variety of data types can lower the requirement to integrate data from multiple data platforms, while data virtualization functionality can help reduce data ingestion overheads by enabling data to be accessed and queried where it resides.

Additionally, databases that are tightly integrated with data management tools that deliver intuitive (and augmented) data exploration functionality can help to lower data preparation overhead and provide the potential to accelerate all aspects of the AI lifecycle. AI-enabled databases also have a significant role to play in lowering the amount of time spent on development and training, as well as testing, deployment and inference – for instance, through the integration of development languages and frameworks, such as Python, GO, JSON and Jupyter notebooks.

An AI-enabled database can also include machine learning capabilities embedded directly within the database. Advantages include having fewer systems to manage through the avoidance of stand-alone environments for machine learning development, as well as enabling access to the complete data, as opposed to having to down-sample and pull data out of the database and move it to a separate system. Another related benefit is the reduction of data movement overhead.

While the strategy and approach vary among the various database vendors in the market, in principle, the idea is the same: to drive efficiencies between what have historically been two different activities, carried out in different environments.

Conclusions

At a high level, an AI-enabled database incorporates AI in a couple of ways: through embedding AI to improve data management performance and using AI functionality to accelerate the development of machine learning models and deployment of AI applications. Embedding AI as part of the functionality of the data management system could include built-in capabilities that enable query optimization and workload management, for example. But AI can also be fused with security, availability, scaling, tuning, infrastructure failure, backup and recovery, and many other tasks that are administered by DBAs.

However, data engineers and DBAs should not necessarily be threatened by AI functionality because primarily, AI has the ability to handle many administrative duties that might otherwise be routine or repetitive. In that sense, database administrators and architects can focus on higher-impact tasks where AI is less mature, such as architecture planning, data modeling, lifecycle management, application-specific tuning, service level management, test data management and database sharding. The important point is that while many database management tasks can be automated with AI, there are many that still require human expertise.

AI functionality can also be used as part of the data management process to accelerate data access, the development of machine learning models, and the deployment of AI applications. This approach has its benefits because it can serve several different personas. For instance, data and business analysts may not have SQL skills to query the data sufficiently, so AI-based functionality may facilitate the searching of the data with natural language processing such that the query syntax is abstracted from the user.

Additionally, built-in AI functionality can enable application developers, particularly if the database system enables integration with developer languages, tools and frameworks. For data scientists, AI functionality may include machine learning algorithms and functions that can be used to better understand, prepare and evaluate data.

Even though the notion of a fully AI-enabled database is still in its early stages, there are a number of benefits that await enterprises with the courage to deploy these systems that go beyond pure database enablement, system performance, data virtualization and reduced administrative burden. The true benefits may be the outcomes of these synergies, such as significant competitive advantages over enterprises that are AI laggards.

Recommendations

- Identify potential use cases for AI as soon as possible, especially tasks are that are overtly manual in nature. Enterprises that are not investing in the development of AI applications risk being left behind by those that are already having success with early initiatives.
- Enterprises must make certain that they are implementing databases designed to support AI functionality by embedding AI at the data level, ensuring that AI enables the full scope of the data management lifecycle, from ingestion to curation and discovery, as well as driving applications that are built on that data.
- Data architects should ensure that they are delivering data infrastructure that is capable of supporting the rapid development of AI applications through direct support for machine learning tools and frameworks and the acceleration of the complete AI data pipeline.
- Enterprises should explore the potential impact of automation on the role of data engineers and DBAs, identifying repetitive tasks that can potentially be automated, as well as the higher-impact tasks that will benefit from human expertise.
- Enterprises need to ensure buy-in from multiple stakeholders and users. While many organizations are open to the idea of AI initiatives, many lack the fundamental knowledge and background to truly understand what it means to manage these types of systems. A few simple successful projects tend to go a long way within organizations to ensure stakeholder agreement.

IBM Hybrid Data Management is the AI optimized foundation for data that fully integrates machine learning and analytics to drive insights for businesses of all sizes and across all industries. It includes operational databases, data warehousing, data lakes and fast data, with options for on-premises, IBM Cloud Pak for Data and public cloud deployments.

With IBM's Hybrid Data Management strategy empowered by AI, organizations can gain operational efficiencies significantly reducing costs through multiple form factors to liberate data into a fully flexible spectrum of data management services while keeping a consistent and simplistic experience. AI powered databases are able to utilize data from all sources with a holistic approach while providing organizations an enriched view of customers in a self-serviced manner.

For more information see:

Modernizing your Information Architecture with AI (webinar)

www.ibm.com/db2

Content provided by IBM

About 451 Research

451 Research is a leading information technology research and advisory company focusing on technology innovation and market disruption. More than 100 analysts and consultants provide essential insight to more than 1,000 client organizations globally through a combination of syndicated research and data, advisory and go-to-market services, and live events. Founded in 2000 and headquartered in New York, 451 Research is a division of The 451 Group.

NEW YORK
Chrysler Building
405 Lexington Avenue,
9th Floor
New York, NY 10174
+1 212 505 3030

SAN FRANCISCO
505 Montgomery Street,
Suite 1052
San Francisco, CA 94111
+1 212 505 3030

LONDON
Paxton House
30, Artillery Lane
London, E1 7LS, UK
+44 (0) 203 929 5700

BOSTON
75-101 Federal Street
Boston, MA 02110
+1 617 598 7200

451 Research®